

DECnet/E
Network Programming in COBOL
Order No. AA-H503B-TC

January 1982

This manual presents procedures for DECnet/E services available to the COBOL programmer, permitting him to design and implement programs that exchange data with other programs running in a DECnet network.

| | | |
|--------------------------------------|-----------------|-------------|
| OPERATING SYSTEM AND VERSION: | RSTS/E | V7.1 |
| SOFTWARE VERSION: | DECnet/E | V2.0 |

To order additional copies of this document, contact your local
Digital Equipment Corporation Sales Office.

digital equipment corporation • maynard, massachusetts

First Printing, January 1982

The information in this document is subject to change without notice and should not be construed as a commitment by Digital Equipment Corporation. Digital Equipment Corporation assumes no responsibility for any errors that may appear in this document.

The software described in this document is furnished under a license and may only be used or copied in accordance with the terms of such license.

No responsibility is assumed for the use or reliability of software on equipment that is not supplied by DIGITAL or its affiliated companies.

Copyright © 1982 by Digital Equipment Corporation

The postage-prepaid READER'S COMMENTS form on the last page of this document requests the user's critical evaluation to assist us in preparing future documentation.

The following are trademarks of Digital Equipment Corporation:

DIGITAL
DEC
PDP
DECUS
UNIBUS
COMPUTER LABS
COMTEX
DDT
DECCOMM
ASSIST-11
VAX
DECnet
DATATRIEVE

DECsystem-10
DECtape
DIBOL
EDUSYSTEM
FLIP CHIP
FOCAL
INDAC
LAB-8
DECSYSTEM-20
RTS-8
IAS
TRAX

MASSBUS
OMNIBUS
OS/8
PHA
RSTS
RSX
TYPESET-8
TYPESET-11
TMS-11
ITPS-10
SBI
PDT

310ALL

Contents

| | Page |
|---|------|
| Preface | vii |
| Part I Introduction | |
| Chapter 1 Introduction | |
| 1.1 Basic Terms and Concepts | 1-1 |
| 1.2 DECnet Networks | 1-2 |
| 1.3 The DIGITAL Network Architecture (DNA) | 1-3 |
| 1.4 Network Routing. | 1-3 |
| 1.5 Overview of DECnet/E Structure | 1-4 |
| Chapter 2 Overview of the Concise and Extended Interfaces | |
| 2.1 The Logical Link | 2-1 |
| 2.2 Network Addressing | 2-3 |
| 2.3 The Concise COBOL Interface | 2-3 |
| 2.4 The Extended COBOL Interface | 2-5 |
| 2.5 Comparison of the Interfaces | 2-7 |
| Part II Network Programming with the Concise COBOL Interface | |
| Chapter 3 Concise COBOL Interface Calls | |
| 3.1 Programming Background | 3-1 |
| 3.1.1 Coding | 3-1 |
| 3.1.2 Compiling and Linking | 3-2 |
| 3.2 Connect Request Call (CNTCON) | 3-3 |
| 3.3 Connect Accept Call (CNTACP) | 3-6 |
| 3.4 Send Record Call (CNTSND) | 3-9 |
| 3.5 Receive Record Call (CNTRCV) | 3-11 |
| 3.6 Disconnect Call (CNTDIS) | 3-13 |

Part III Network Programming with the Extended COBOL Interface

Chapter 4 Extended Discussion of Logical Links

| | | |
|-------|---|-----|
| 4.1 | Creating a Logical Link in the Extended COBOL Interface | 4-1 |
| 4.2 | Advantages of Logical Links | 4-4 |
| 4.2.1 | Efficient Line Usage | 4-4 |
| 4.2.2 | Separate Data by Intended Use | 4-6 |
| 4.2.3 | Interrupt Messages | 4-7 |
| 4.2.4 | Flow Control | 4-7 |
| 4.3 | DECnet/E Logical Link Limitations. | 4-8 |

Chapter 5 Overview of Extended DECnet/E Message Services

| | | |
|-------|---|------|
| 5.1 | General Remarks | 5-2 |
| 5.2 | Registering with the Monitor: Declare Receiver Call | 5-3 |
| 5.2.1 | Declaring Identity | 5-3 |
| 5.2.2 | Declaring Intended Usage | 5-4 |
| 5.2.3 | DECnet/E Use of the Declare Receiver Call | 5-5 |
| 5.3 | Remove Receiver Call | 5-5 |
| 5.4 | Get Local Node Parameters Call | 5-5 |
| 5.5 | Log User Event Call | 5-6 |
| 5.6 | Send/Receive Local Data Message | 5-6 |
| 5.7 | Send/Receive Connect Initiate Message | 5-7 |
| 5.8 | Send/Receive Connect Confirm Message. | 5-8 |
| 5.9 | Send/Receive Connect Reject Message | 5-9 |
| 5.10 | Send/Receive Network Data Message | 5-10 |
| 5.11 | Send/Receive Interrupt Message | 5-10 |
| 5.12 | Send/Receive Link Service Message. | 5-11 |
| 5.13 | Send/Receive Disconnect Message | 5-11 |
| 5.14 | Send/Receive Link Abort Message | 5-12 |

Chapter 6 Network Addressing and Flow Control

| | | |
|-------|---|------|
| 6.1 | Network Addressing | 6-1 |
| 6.1.1 | Declaring Network Names and Object Type Codes | 6-1 |
| 6.1.2 | Handling of Incoming Connect Requests | 6-2 |
| 6.1.3 | Multiple Copies | 6-3 |
| 6.1.4 | Automatic Job Startup | 6-4 |
| 6.1.5 | Examples of Network Addressing | 6-6 |
| 6.2 | Flow Control | 6-9 |
| 6.2.1 | Transmit Queue Management. | 6-10 |
| 6.2.2 | Backpressure Flow Control | 6-13 |
| 6.2.3 | Program-Regulated Flow Control | 6-15 |
| 6.2.4 | Interrupt Message Flow Control | 6-19 |
| 6.2.5 | Programming Hints for Flow Control. | 6-20 |

Chapter 7 Extended COBOL Interface Call Formats

| | | |
|-------|---|------|
| 7.1 | Programming Background | 7-1 |
| 7.1.1 | Coding | 7-1 |
| 7.1.2 | Compiling and Linking | 7-5 |
| 7.2 | Declare Receiver (MDCL) | 7-6 |
| 7.3 | Remove Receiver (MREM) | 7-12 |
| 7.4 | Get Local Node Parameters (NTLN) | 7-14 |
| 7.5 | Log User Event (NTEV) | 7-17 |
| 7.6 | Send Calls. | 7-21 |
| 7.6.1 | Send Local Data Message (MSLD) | 7-21 |
| 7.6.2 | Send Connect Initiate Message (NTCI) | 7-24 |
| 7.6.3 | Send Connect Confirm Message (NTCC) | 7-32 |
| 7.6.4 | Send Connect Reject Message (NTCR) | 7-36 |
| 7.6.5 | Send Network Data Message (NTDM) | 7-39 |
| 7.6.6 | Send Interrupt Message (NTIN) | 7-44 |
| 7.6.7 | Send Link Service Message (NTLS) | 7-48 |
| 7.6.8 | Send Disconnect Message (NTDI) | 7-52 |
| 7.6.9 | Send Link Abort Message (NTLA) | 7-55 |
| 7.7 | Receive (MRCV) | 7-58 |
| 7.7.1 | Format of Received Local Data Message | 7-64 |
| 7.7.2 | Format of Received Connect Initiate Message | 7-65 |
| 7.7.3 | Format of Received Connect Confirm Message | 7-69 |
| 7.7.4 | Format of Received Connect Reject Message | 7-70 |
| 7.7.5 | Format of Received Network Data Message | 7-71 |
| 7.7.6 | Format of Received Interrupt Message | 7-72 |
| 7.7.7 | Format of Received Link Service Message | 7-73 |
| 7.7.8 | Format of Received Disconnect Message | 7-75 |
| 7.7.9 | Format of Received Link Abort Message | 7-76 |

Appendix A Object Type Codes

Appendix B NSP Reason Codes for Connect Reject and Link Abort

Index

Figures

| | | |
|-----|--|-----|
| 1-1 | Computer Network Composed of Five Nodes | 1-2 |
| 2-1 | NSP Handles Multiplexing and Demultiplexing of Data over Physical Lines | 2-2 |
| 2-2 | Establishing a Logical Link | 2-2 |
| 4-1 | Logical Link Connections. | 4-3 |
| 4-2 | Logical Links Provide Line Efficiency. | 4-5 |

| | | |
|-----|--|------|
| 4-3 | User Link Addresses, Remote Link Addresses, and Local Link Addresses for Logical Links. | 4-5 |
| 4-4 | Data Is Queued for Logical Links in Separate Streams. | 4-6 |
| 4-5 | Queuing of Interrupt Messages for Logical Links. | 4-7 |
| 6-1 | Network Addressing Used by NFT and FAL. | 6-7 |
| 6-2 | Network Addressing Example — Different Declare Receiver for Local and Remote Start. | 6-8 |
| 6-3 | Network Addressing Example — Declaring Identity by Name Alone . . . | 6-9 |
| 6-4 | How Transmit Queue Flow Control Affects a DECnet/E Program. . . . | 6-12 |
| 6-5 | How Backpressure Flow Control Affects a DECnet/E Program as a Transmitter. | 6-14 |
| 6-6 | How the Flow Control Option Selected by the Remote Program Affects a DECnet/E Program as a Transmitter | 6-18 |
| 6-7 | How Interrupt Message Flow Control Affects a DECnet/E Program as a Receiver | 6-20 |
| 7-1 | Format of Connect Data Block on Send Connect Initiate Message . . . | 7-26 |
| 7-2 | Format of Received Connect Data Block | 7-66 |

Tables

| | | |
|-----|---|------|
| 5-1 | Summary of Calls for Interprogram Communication | 5-1 |
| 7-1 | Types of Receiver Access | 7-7 |
| 7-2 | System Validation of Name, Object, and Access Parameters | 7-9 |
| 7-3 | Format of Connect Data Block on Send Connect Initiate Message . . . | 7-26 |
| 7-4 | Sender Selection Summary. | 7-60 |
| 7-5 | Format of Received Connect Data Block | 7-66 |

Preface

Manual Objectives

This manual describes the DECnet/E services available to a COBOL programmer, permitting him to code programs that exchange data with other programs running in a DECnet network.

Intended Audience

This manual is written for the experienced programmer. While the reader is not expected to be knowledgeable of network programming, he is expected to be completely familiar with both the COBOL programming language and the general mechanisms for interfacing with the RSTS/E operating system.

Prerequisite Reading

The reader is expected to be completely familiar with the concepts and features of the COBOL programming language. The *RSTS/E COBOL User's Guide* tells how to compile COBOL programs in the RSTS/E system. Also, the *RSTS/E Task Builder Manual* describes TKB, the linker used for building RSTS/E tasks from compiled COBOL programs.

The reader is also expected to be familiar with the concepts and features available for interfacing to the RSTS/E system monitor, as presented in the following manuals:

RSTS/E System User's Guide
RSTS/E Programming Manual

Related Documents

To obtain a general understanding of the DECnet environment — internal and external — the reader is directed to the *Introduction to DECnet*.

DECnet/E also provides utility programs that allow a terminal user to communicate with another terminal user, to manipulate files across the network, and to copy the contents of one storage device to a device on another system. The information needed to run and use these utilities is given in the companion manual, *DECnet/E Guide to User Utilities*.

Two other companion manuals tell how to generate and start a DECnet/E system (the *DECnet/E System Installation Guide*), control and manage a running DECnet/E system (the *DECnet/E System Manager's Guide*), and give other information useful to a system manager.

Three more DECnet/E manuals are available describing the DECnet/E message send/receive services for four other languages:

DECnet/E Network Programming in MACRO-11

DECnet/E Network Programming in FORTRAN

DECnet/E Network Programming in BASIC-PLUS and BASIC-PLUS-2

The more advanced network programmer may want to refer to the following two manuals for information concerning various network management functions:

DIGITAL Network Architecture, Network Services Functional Specification

DIGITAL Network Architecture, Network Management Functional Specification

Structure of This Manual

This manual is divided into three main parts. Part I, "Introduction," presents a general discussion of the DECnet/E environment (Chapter 1), and an overview of the two separate interfaces available to the COBOL programmer for accessing the DECnet/E services (Chapter 2).

Part II, "Network Programming with the Concise COBOL Interface," describes the formats of the calls available to COBOL programmers who choose to use this abbreviated interface (Chapter 3).

Part III, "Network Programming with the Extended COBOL Interface," presents an expanded tutorial discussion of the DECnet/E features available for sending and receiving messages (Chapters 4 through 6), and detailed descriptions of the extended interface calls used to access those features (Chapter 7).

Part I

Introduction



Chapter 1

Introduction

DECnet/E is a combination of hardware and software that extends the capabilities of a RSTS/E system running on a DIGITAL PDP-11 computer to allow a programmer to develop and execute programs that exchange data with remote programs running on another DECnet system in a network.

NOTE

Throughout this manual, the term *remote program* refers to any other program with which a program is communicating using the DECnet facilities. Such programs need not be located on another node in the network. It is possible to use DECnet/E to communicate with other programs at the local RSTS/E node. During development of network application programs, the whole system can be debugged locally and later distributed to remote nodes as required.

This manual describes the features available to network programmers in the COBOL programming language on RSTS/E systems.

1.1 Basic Terms and Concepts

The term *network* is used here to refer to two or more computer systems connected such that they can exchange information. The computer systems, called *nodes*, are connected by communications paths called *physical links*. Physical links can be established through normal (switched) telephone circuits, several varieties of leased (private) lines, coaxial cables, or even satellite transmission facilities provided by several carriers.

In the network shown in Figure 1-1, a user at node BOSTON would refer to BOSTON as the *local node* and to nodes DENVER, DALLAS, LONDON, and LA as *remote nodes*. At node BOSTON, the nodes DENVER, DALLAS, and LONDON are physically *adjacent nodes*. That is, there is a direct physical link to these nodes with no intervening nodes. To a user at node LA, DENVER is the only adjacent node.

A *path* is the route over which data travels from its source to its destination within the network. *Path length* is the distance from the source node to the destination node, measured in hops. A *hop* is equal to a physical line between two adjacent nodes. In Figure 1-1, the path length between nodes DALLAS and BOSTON is one hop and between LA and LONDON is three hops. The *network diameter* is the maximum path length between any two nodes. The network shown in Figure 1-1 has a network diameter of three hops.

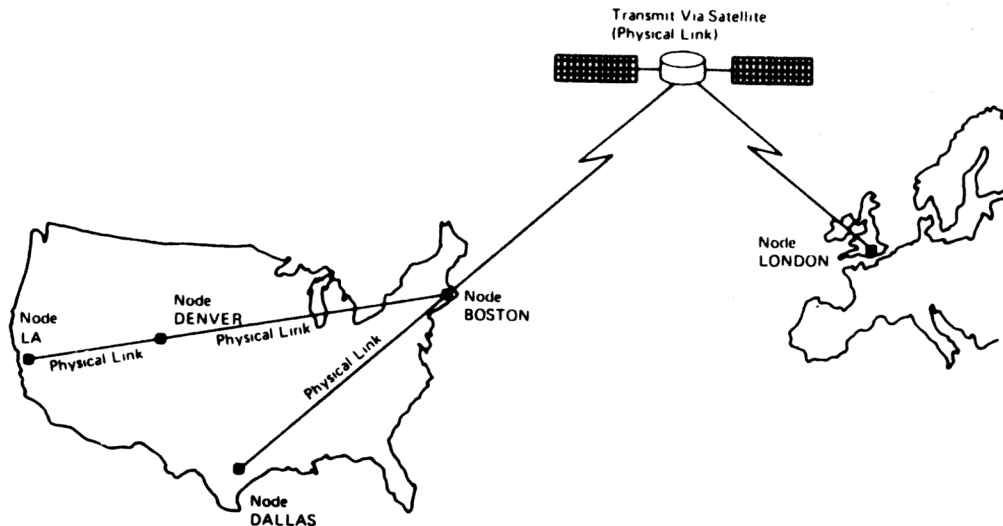


Figure 1-1: Computer Network Composed of Five Nodes

1.2 DECnet Networks

DECnet as a whole refers to the hardware and software that allows different DIGITAL systems to be connected to form networks. DECnet has been designed to allow each system to function as a separate entity but still allow access to resources that are distributed throughout the various systems in the network. A DECnet/E system provides all the capabilities of a normal RSTS/E system as well as the networking capabilities described in this manual and in the companion DECnet/E manuals listed in the Preface.

All DECnet implementations are based on a design structure called the DIGITAL Network Architecture (DNA). This architecture allows network implementations on the various DIGITAL systems to evolve (to provide more and more capability as time goes on) within a basic common structure. The structure ensures that implementations providing fewer features will always work with implementations providing more.

Perhaps a more immediate aspect of this process of evolution, however, is that at any given time capabilities can indeed differ from system to system. We do not attempt to describe here the DECnet implementations for all the different DIGITAL systems. Nevertheless, to use DECnet/E, you do need to give some thought to the other DIGITAL systems in your network.

If your network consists entirely of DECnet/E Version 2.0 nodes, this document provides all the information necessary for you to use DECnet/E to send and receive messages in COBOL programs. If your network contains other DECnet systems, however, you should check the DECnet documentation for those systems to be sure that the corresponding capabilities exist. For example, you can design network programs to be started automatically on RSTS/E systems. This feature is not available on all DECnet systems, however, so your design should not necessarily assume this capability.

A brief overview of the functional capabilities of the various DECnet implementations can be found in *Introduction to DECnet*.

1.3 The DIGITAL Network Architecture (DNA)

DECnet/E is implemented according to a set of rules, or protocols, governing the format, control, and sequencing of data exchange from the user program level down through the physical link level. These protocols are defined by the DIGITAL Network Architecture (DNA) — a logical structure that provides the model for all DECnet implementations.

DNA consists of several layers, each defining a distinct set of network functions and a set of rules for implementing those functions. Each DECnet implementation consists of software modules that perform these DNA-defined functions according to DNA-defined protocols. The relevant protocols for network programming are called the Network Services Protocol (NSP), the Transport protocol, and the Digital Data Communications Message Protocol (DDCMP).

The Network Services Protocol defines the rules for communication between programs in a network over paths called logical links. Logical links, described in detail in Chapter 4, allow many different data streams to be multiplexed onto the physical link for transmission and separated at the receiving node for delivery to the appropriate program.

The Transport protocol defines the rules for determining the actual physical path, or *route*, along which data travels to its destination.

Physical link control is achieved in DECnet by implementation of the Digital Data Communications Message Protocol (DDCMP). DDCMP ensures an error-free, sequential data path over a generally error-prone medium. This is accomplished with the Cyclic Redundancy Check (CRC) for error detection, retransmission for error corrections, and numbered data segments to ensure sequential transmission of data.

1.4 Network Routing

DECnet/E Version 2.0 is a Phase III DECnet product. (Refer to *Introduction to DECnet* for a discussion of the differences between Phase II and Phase III DECnet products.) As such, it supports the adaptive routing feature of Phase III, providing the user with the capability of communicating with other nodes in the network regardless of whether or not they are directly connected to the local node.

Networks can consist of the following three types of nodes:

- **Routing nodes** — Phase III nodes connected to multiple communications lines, supporting route-through capabilities.
- **Nonrouting nodes (end nodes)** — Phase III nodes connected to a single communications line.
- **Phase II nodes** — nodes running a previous generation of the DECnet architecture.

Phase III nodes can communicate with and are compatible with Phase II nodes. However, Phase II nodes do not acquire any new capabilities by being connected to Phase III nodes, and the restriction that Phase II nodes can only communicate with adjacent nodes continues to apply. Thus, a Phase III node can communicate with any other Phase III node, *providing the path goes through Phase III nodes only*.

The adaptive routing feature is implemented through an algorithm that chooses the routing path with the lowest associated cost. Cost is computed as the sum of the costs of the lines over which the message is transmitted. The individual line cost parameters are assigned by the system manager and input into the system with the Network Control Program. (See the *DECnet/E System Manager's Guide* for a full discussion of line costs.) The algorithm automatically adjusts the routing when network topology or line cost changes occur.

1.5 Overview of DECnet/E Structure

The DECnet/E features for network communication are implemented as part of the RSTS/E system monitor. The manner in which these features are accessed differs with the type of user.

User programs interface with DECnet/E through NSP. Although it is not a separate entity or program, the term NSP is used throughout this manual to refer to the software that allows a user program to establish and exchange data over logical links. In other words, NSP is the implementation of the Network Services Protocol. The RSTS/E COBOL programmer accesses NSP services by coding calls to MACRO subprograms. NSP services can also be called from MACRO-11, BASIC-PLUS, BASIC-PLUS-2, or FORTRAN programs, as described in associated DECnet/E Network Programming Manuals.

Terminal users access certain network capabilities through various DECnet/E utility programs. The utilities TLK (Talk) and LSN (Listen) allow two terminal users to type messages to each other. The utilities NFT (Network File Transfer) and FAL (File Access Listener) work together to allow network file manipulation. They are implemented according to the Data Access Protocol (DAP) so that files can be exchanged with other DECnet systems having the same capabilities. The NET utility allows the terminal user to log into and use a remote RSTS/E system as though the local terminal were directly connected to the remote system, and NETCPY copies entire devices between RSTS/E nodes. (See the *DECnet/E Guide to User Utilities* for details on these utility programs.)

The system manager uses the features provided by the Network Control Program (NCP) to monitor and control a node running in a network. (See the *DECnet/E System Manager's Guide* for details on running NCP.)

Some elements of DECnet/E are not directly accessible to any user. These lower level elements of DECnet/E include the network Transport module (TRN) that implements the Transport protocol, several hardware devices (the DMC11, DMR11, DMV11, and DMP11), and their associated software device drivers.

The device drivers are responsible for handling messages received from and sent over the physical communication lines. NSP gives a message to be sent to Transport, which selects an appropriate data path based on the destination of the message. Transport then passes the message to the proper device driver for actual transmission. The drivers manipulate the device registers to cause message transmission to occur.

To handle incoming messages, the drivers include buffer management functions that allow messages to be received from the physical links. When a message is received, it is passed to Transport, which checks the destination of the message. If the message is for the local node, Transport passes it to NSP for further processing. If the message is destined for another node in the network, Transport selects an outgoing data path and passes the message to one of the drivers for transmission.

The hardware devices are intelligent communication controllers that connect the PDP-11 to physical communications lines. The DMC11 and the DMR11 control lines that connect to only one other system in the network (point-to-point lines) while the DMV11 and the DMP11 control lines that may connect to more than one other system (multipoint lines). The physical lines can include cables, modems and telephone circuits, or even satellite transmission facilities. Each controller contains a microprocessor, its own memory, and microcode that implements the DDCMP protocol. The implementation of DDCMP in firmware considerably reduces the software overhead for the DECnet/E system.

Chapter 2

Overview of the Concise and Extended Interfaces

The COBOL programmer can access the DECnet/E message services in one of two forms: the concise COBOL interface or the extended COBOL interface. This chapter introduces the logical link and network addressing, fundamental to all network programming and briefly describes the concise and extended interfaces. To help you decide which interface to use, the concluding section summarizes the advantages and disadvantages of each interface.

2.1 The Logical Link

A logical link is a software path for the exchange of data between two programs running in a network.

In a network of DECnet nodes, many programs use the same physical path to exchange data. Data being sent from one program to another over a logical link is interspersed (multiplexed) with data sent over other logical links, transmitted over the same physical path, and separated (demultiplexed) at the receiving end. The Network Services Protocol (NSP) handles this multiplexing and demultiplexing. It accepts outgoing data from local programs and formats it for transmission by the communications hardware. It also accepts incoming data from remote nodes and separates it into individual logical link streams that are then delivered to the appropriate local programs (see Figure 2-1).




Figure 2-1: NSP Handles Multiplexing and Demultiplexing of Data over Physical Lines

Two programs must first agree to communicate before a logical link is established (see Figure 2-2). Once the link is established, both programs can send and receive data over the logical link. When transmission is finished, either program can disconnect the link.

This process — establishing a logical link, transmitting and receiving data, and disconnecting the link — is basically the same for all implementations of DECnet. With the concise COBOL interface, your program can establish one logical link with another program in the network. Programs using the extended COBOL interface can establish several logical links with the same or different network programs. An expanded discussion of logical links, directed at the extended COBOL interface user, is given in Chapter 4.




Figure 2-2: Establishing a Logical Link

2.2 Network Addressing

A DECnet/E program can declare its identity in the network with an object type code, a name, or both. In the concise COBOL interface, only the program that accepts a logical link connection needs to declare its identity. The requestor identifies the program with which it wishes to connect.

A network program's name can be from one to six characters in length. Valid characters are uppercase, alphanumerics and the special characters "\$" (dollar sign), "." (period), and "_" (underline). Embedded and leading spaces are invalid. The name must be unique to the node. That is, no other program at the local node can declare its identity with the same network name at the same time.

An object type code is a number in the range of 1 to 255. Object type codes simplify the identification of programs performing identical functions in the network. For example, all DECnet TLK utilities, that allow terminal users to communicate with each other in the network, declare their identity with an object type code of 16. The range of object type codes from 1 to 127 is reserved for DECnet use. The numbers ranging from 128 through 255 can be used by a network installation for its own commonly used network programs. Currently assigned object type codes for DECnet use are given in Appendix A.

DECnet/E allows more than one program to declare their identity with the same object type code at the same time. Thus, multiple copies of a program that declares its identity with an object type code only, or an object type code and a unique name, can be executed by terminal users or through the RSTS/E BATCH processor.

Programs can also be set up under DECnet/E to be started automatically to meet incoming requests for logical links. The system manager can define such programs, relating a name or object type code to the name of a file on disk. A program defined in this fashion will be loaded from disk and started automatically by NSP when a logical link is requested with the defined name or object type code.

An expanded discussion of network addressing, with features available to those using the extended COBOL interface, is given in Chapter 6. The *DECnet/E System Manager's Guide* describes how programs are installed for automatic startup. At this point, you should simply be aware that multiple copies of your program can be executed at a RSTS/E node if you declare the program's identity with an object type code and that object type codes are intended for use in a consistent fashion throughout the network. You would not, for example, use an object type code of 129 for a payroll program at node A and for an inventory program at node B.

2.3 The Concise COBOL Interface

The design of the concise COBOL interface allows the COBOL programmer to regard the network as a *file*. That is, you can *open* the file (establish a logical link), send or receive *records* (up to 32767 bytes of data in any COBOL-definable category), and *close* the file (disconnect the logical link). Only one

network file can be open at a time. That is, you can establish only one logical link with another program in the network. Once established, both programs can send and receive data over the logical link and either program can disconnect the logical link.

The following five calls are available with the concise COBOL interface:

1. **Connect Request Call (CNTCON):** The CNTCON call is used to request a logical link connection with another program in the network. The call identifies the node and the remote program's name or object type code. The request is forwarded to the specified node. If a program is running there that has already issued a Connect Accept call, identifying itself with the proper object type code or name, or if a program with the proper identity has been defined for automatic startup, the request for a connection is delivered. Information can be included to permit the remote system or program to verify the calling program's right to use the remote system or program's services. If the remote program does not accept the connection within five minutes, the CNTCON call terminates with an error value returned to an argument specified in the call.
2. **Connect Accept Call (CNTACP):** The CNTACP call is used to accept a logical link connection request from another program in the network. In the call, the program identifies itself with a network name or object type code. The calling program can also supply information to be used by NSP to check the requestor's rights to access the program's services. If this information does not match that specified by the remote program in its connect request, the request for the logical link is rejected.
3. **Send Record Call (CNTSND):** The CNTSND call allows a program to send up to 32767 bytes of data to a remote program over an established logical link. As with all RSTS/E input and output, the CNTSND call is synchronous. That is, control returns to the calling program when all the data is sent and has been received by the remote system.
4. **Receive Record Call (CNTRCV):** The CNTRCV call delivers up to 32767 bytes of data to a data area specified by the calling program. Like CNTSND, the CNTRCV call is synchronous.
5. **Disconnect Call (CNTDIS):** The CNTDIS call is used to disconnect the established logical link. Either program can disconnect a logical link. A disconnection will cause a CNTRCV issued by the other end of the logical link to terminate with an end-of-file indication.

With the concise COBOL interface, you should design your programs so that two possible deadlock situations cannot occur: simultaneous send calls or simultaneous receive calls at both ends of a logical link.

If both programs using a logical link try to send large amounts of data at the same time, and neither program's allotted portion of system buffer space is large enough to hold its incoming record, both programs will suspend execution indefinitely, waiting for the I/O operation to complete. Neither program can receive, thereby alleviating the congestion, because neither send can complete.

(No send deadlock will occur if the maximum length of the records being sent is 183 bytes. Longer records can be broken into multiple records, 183 bytes or less, and then be reconstructed at the receiving end.)

Similarly, if both programs using a logical link execute a CNTRCV at the same time, without sending data first, both will wait indefinitely for data to be received.

In the case where one program only sends data and the other program only receives data, neither of the deadlock situations can ever occur. In the more general case, where both programs can send and receive, a message dialog should be established between the programs so that each program recognizes its turn to send or receive.

2.4 The Extended COBOL Interface

The extended COBOL interface allows more finely-tuned network operations than the concise COBOL interface. A program can establish more than one logical link with the same or separate remote programs. Also, a program can select flow control options to regulate the flow of incoming data. Such a program can be designed so that it will never wait indefinitely after trying to send or receive.

The calls were not designed, however, with COBOL applications in mind. Data is sent and received in units called *segments*. The maximum segment size is determined by the communications hardware, with the largest possible segment being 532 bytes. The flow control options mentioned previously regulate flow in terms of segments or *logical messages* (one or more segments, as defined by a bit in the send call argument). To realize the full power of the extended interface, you should be prepared to access one-byte binary items and set bits — not really easy with COBOL — and even code MACRO-11 (assembly language) subprograms to execute RSTS/E monitor directives. The extended COBOL interface calls reflect the way in which the network message services are implemented in the RSTS/E monitor. The same functions are available to MACRO-11, FORTRAN, BASIC-PLUS, and BASIC-PLUS-2 users on DECnet/E systems.

In any case, 14 calls are available. The following discussion will give you some idea of the scope of network message services available with the extended COBOL interface. A full discussion of the interface is given in Part III.

1. **Declare Receiver Call (MDCL):** The MDCL call allows a network program to register for network operations. The program identifies itself by object type code, name, or both. It defines the maximum number of logical links it will handle, the maximum number of messages to be queued for it in system buffer space, and a limit on system buffer space to be used on its behalf.
2. **Remove Receiver Call (MREM):** The MREM call terminates message operations, destroying any logical links still active.
3. **Get Local Node Parameters Call (NTLN):** The NTLN call returns information to the calling program concerning the local node's network parameters.

4. **Log User Event Call (NTEV):** The NTEV call permits a user-written program to queue an event to the system event processor for logging.
5. **Send Local Data Message (MSLD):** Programs using the extended COBOL interface can send up to 532 bytes of data to a program at the local node with this call. It is not necessary to establish a logical link in this case.
6. **Send Connect Initiate Message (NTCI):** The NTCI call is used to request a logical link with another program in the network. A 120-byte Connect Data Block is used to address the remote program. It includes numeric values stored as one-byte binary items. In addition, the program specifies a link identifier (since it can have more than one link), the flow control option desired, and a maximum segment size for receiving data.
7. **Send Connect Confirm Message (NTCC):** The NTCC call is used to accept a request for a logical link connection. As with the NTCI call, the program specifies a link identifier, the type of flow control, and a maximum segment size for receiving data.
8. **Send Connect Reject Message (NTCR):** The NTCR call is used to reject a request for a logical link connection.
9. **Send Network Data Message (NTDM):** The NTDM call is used to transmit user data to a remote program over an established logical link. The amount of data sent is limited to the maximum the remote program wishes to receive. Sending a Network Data Message is subject to flow control imposed by the local system, the remote system, and the remote program. A full discussion of flow control is given in Chapter 6.
10. **Send Interrupt Message (NTIN):** The NTIN call is used to transmit high-priority data to a remote program over an established logical link. All DECnet systems maintain separate data streams for Interrupt and Data Messages. The Interrupt Message is delivered to the remote program ahead of all Data Messages waiting to be processed at the remote end.
11. **Send Link Service Message (NTLS):** The NTLS call is used to request data over a logical link when the program has requested either segment or logical message flow control. A full discussion of these types of flow control is given in Chapter 6.
12. **Send Disconnect Message (NTDI):** The NTDI call is used to break an established logical link provided that all Network Data, Interrupt, and Link Service Messages previously sent over the link have been acknowledged by the remote NSP. If they have not been acknowledged, the call will terminate with an error and the link will not be broken.

13. **Send Link Abort Message (NTLA):** The NTLA call is used to break a logical link immediately. No attempt is made to ensure that messages previously sent over the link have been acknowledged by the remote system.
14. **Receive Call (MRCV):** The MRCV call is used to retrieve a pending message from system buffer space and deliver it to an area defined within the user program. Any of nine message types can be received, corresponding to the nine message types in items 4 - 12.

2.5 Comparison of the Interfaces

In summary, the concise COBOL interface is much simpler to use than the extended COBOL interface. Your program can establish one logical link with another program in the network and send or receive data over the link with calls similar to file I/O operations. The CNTCON and CNTACP calls are similar to opening a file. Your program issues one or the other, depending on whether it is requesting or accepting a logical link. The CNTSND and CNTRCV calls are similar to write and read operations — you send data or receive data across the network. Likewise, the CNTDIS call corresponds to closing a file.

If, however, you want to take advantage of the full power of DECnet/E message services, the extended COBOL interface is there for you to use. The extra features available with the extended interface allow you to establish multiple logical links, send Interrupt Messages or Local Data Messages, and closely control the amount of data sent to your program.

Part II
Network Programming with the
Concise COBOL Interface

Chapter 3

Concise COBOL Interface Calls

This chapter presents the specific calls, with detailed formats, used to access the DECnet/E message services with the concise COBOL interface.

3.1 Programming Background

To use the DECnet/E message services available with the concise COBOL interface, you code subprogram calls with the CALL statement, including the USING phrase to pass operand values to the subprogram. These subprogram calls reference MACRO-11 subprograms stored in the system library that must be linked with your program. Thus, the process involves coding, compiling, and linking. These aspects of using the DECnet/E concise COBOL interface are described briefly in the following subsections. Specific call formats are given in Sections 3.2 - 3.6.

3.1.1 Coding

The general format of the concise COBOL message calls is as follows:

CALL "*name*" USING *arg1*, *arg2*,

where

name is the name of one of the five subprograms available for message services. Each subprogram processes a call at execution time by executing the following steps:

1. Checking for the proper number of arguments in the call.
2. Translating the call and arguments into the form recognizable to the RSTS/E monitor.
3. Passing control to the RSTS/E monitor and, therefore, to NSP. (In DECnet/E systems, NSP is part of the RSTS/E monitor.)

NSP processes the request and returns control to the user program. The status of each call (whether or not an error occurred) is returned as a COMP value to a data item defined as an argument in each call.

arg... is the argument list for the particular call. You must define each argument as a data item in the Data Division of the program.

Note that the message subprograms are MACRO-11 (assembly language) object code subprograms and the numeric arguments passed and returned require definition in the Data Division of your COBOL program such that they fit in one computer word. Such a COMP data item can be defined with PICTURE statements ranging from

PIC 9(1) COMP through PIC 9(4) COMP (for unsigned values)

or

PIC S9(1) COMP through PIC S9(4) COMP (for signed values)

Once a COMP data item is defined as PIC S9(4), for example, values ranging from -32768 through 32767 can be returned to the data item (even though 5 character positions are used). To pass a value of more than four characters, you MOVE (in the Procedure Division) the value to the data item used as an argument in the call, rather than use the VALUE IS clause in the Data Division.

3.1.2 Compiling and Linking

The DECnet/E subprograms are stored as object modules in the RSTS/E system library (LB:SYSLIB.OLB). References to the subprograms will be resolved automatically during the task building process unless you give specific directions otherwise. Follow your normal procedure for compiling, merging ODLs, and task building, as described in the *PDP-11 COBOL User's Guide*.

3.2 Connect Request Call (CNTCON) — Privileged

The CNTCON call is used to request a logical link with another program in the network. The other program is identified, and the calling program's access rights to the remote system's or program's services are specified with a project-programmer number, a password, and an account.

COBOL Call Format:

CALL "CNTCON" USING *stat*, *rnode*, *rname*, *ppn*, *pwd*, *act*, *reason*

Argument Descriptions:

stat

The *stat* argument is a one-word COMP data item (see Section 3.1.1). It is set to zero if the call completes successfully. If an error occurs, *stat* is set to one of the error codes listed under "Returned Status" at the end of this section.

rnode

The *rnode* argument defines the remote node to which the connect request is directed. It consists of one to six uppercase, alphanumeric characters, left-justified in a 6-byte field with space fill. (This happens automatically with the VALUE IS clause or the MOVE statement.) Leading or embedded spaces are invalid. The name must contain at least one alphabetic character. If no node name is to be specified (that is, if the connect request is directed to a program at the local node), the *rnode* argument should be loaded with 6 spaces.

rname

The *rname* argument is a 21-character alphanumeric data item identifying the remote program with which the connection is to be made. Two types of identification are allowed: object type code or name.

An object type code is a number ranging from 1 to 255. A name is an alphanumeric identifier. The call must specify whichever identifier is used by the remote program. Those network addressing features that are of interest to concise COBOL interface users are discussed in Section 2.2.

The *rname* argument can take one of the following two forms:

TASK=*name*

or

objtyp=

where *name* is the alphanumeric name used by the remote program to identify itself as a DECnet program. It can be from 1 to 16 characters long. (Names on DECnet/E systems can be from one to six uppercase, alphanumeric characters or the special characters "\$" (dollar sign), "." (period), and "_" (underline). Embedded or leading spaces are invalid.

The object type code (*objtyp*), if used, can range from 1 to 255. It must be followed by an equal sign (=) character.

The characters used must be left-justified in the 21-character data item, with space fill. This happens automatically with the VALUE IS clause or MOVE statement.

ppn

The *ppn* argument is a 16-character alphanumeric data item. It is a project-programmer number (referred to on some systems as a user identification code, or UIC). The *ppn* is one of three arguments that can be used to define the calling program's access rights to the remote system's or remote program's services, if the remote system or program checks such access. A receiving DECnet/E system simply passes the information on to the receiving program. DECnet/E programs using the concise COBOL interface, however, can specify a corresponding *ppn* argument to be checked against that of the incoming request.

The project-programmer number must be left-justified in the 16-character field, with space fill. This happens automatically with the VALUE IS clause or the MOVE statement.

pwd

The *pwd* argument is an 8-character alphanumeric data item. It is a password — one of three arguments that define this program's rights to the remote system's or the remote program's services, if the remote system or program checks such access. Receiving DECnet/E systems simply pass the information on to the receiving program. DECnet/E programs using the concise COBOL interface, however, can specify a corresponding *pwd* argument to be checked against that of the incoming request.

The password must be left-justified in the 8-character data item, with space fill. This happens automatically with the VALUE IS clause or the MOVE statement.

act

The *act* argument is a 16-character alphanumeric data item. It is an account number — one of three arguments defining this program's access rights to the remote system's or the remote program's services, if the remote system or program checks such access. Receiving DECnet/E systems simply pass the information on to the receiving program. DECnet/E programs using the concise COBOL interface, however, can specify a corresponding *pwd* argument to be checked against that of the incoming request.

The account number must be left-justified in the data item, with space fill. This happens automatically with the VALUE IS clause or the MOVE statement.

reason

The *reason* argument is a one-word COMP data item. It is set to a reason code that further defines the reason for a rejected link request when the *stat* argument is 5 or 6. Reason codes are listed in Appendix B.

Returned Status:

| stat | Meaning |
|------|--|
| 0 | The CNTCON call completed successfully. |
| 1 | The program already has an existing logical link. Only one link at a time can be used with the concise COBOL interface. The current link must be disconnected before a CNTCON call can succeed. |
| 2 | The format of the call is incorrect. For example, the wrong number of arguments was specified in the call or their lengths are not in the proper range. |
| 3 | The CNTCON call was issued by a nonprivileged program. A RSTS/E program must be privileged when it executes a CNTCON call. |
| 4 | The node specified is unknown or is currently inaccessible (the node is shut down, the line is down, and so forth). Or, the node name is invalid. |
| 5 | The CNTCON call was rejected by the remote system or the remote program; or five minutes has elapsed with no response from the remote program. |
| 6 | The physical line to the remote node went down before the remote program responded to the connect request. |
| 14 | The CNTCON call failed for some reason that will require human intervention to fix. For example, NSP has not been enabled and DECnet is not working at the local node; or the program (job) has been removed as a receiver (denied the use of DECnet) by some privileged user on the system. |
| 16 | The format of the <i>rname</i> argument was invalid. |

Example:

The following example shows a request for a logical link connection with a remote program named INVTRY at node BOSTON. A project-programmer number of 1,222 and a password of MYNAME are used. The account is 34C.

(Data Division)

```
01      STAT PIC 9(4) COMP.
01      TASK-NAME PIC X(21) VALUE IS "TASK=INVTRY".
01      NODE-NAME PIC X(16) VALUE IS "BOSTON".
01      PPN PIC X(16) VALUE IS "1,222".
01      PASWRD PIC X(8) VALUE IS "MYNAME".
01      ACCT PIC X(16) VALUE IS "34C".
01      RSN PIC 9(4) COMP.
```

·
·
·

(Procedure Division)

```
CALL "CNTCON" USING STAT, NODE-NAME, TASK-NAME, PPN,  
PASWRD, ACCT, RSN.
```

3.3 Connect Accept Call (CNTACP) — Privileged

The CNTACP call is used to accept a logical link requested by another program in the network. The program issuing the CNTACP call identifies itself either by object type code, name, or both. The project-programmer number, password, and account number arguments specified in the CNTACP call are checked against corresponding values sent by the remote program in the incoming connect request. If they do not agree, the connection will be rejected and the CNTACP call will terminate with a status value of 5 in the *stat* argument.

If no request for a logical link has been received when the CNTACP call is issued, the program will wait until a connection is requested. Thus, a program can be designed to be run manually with the RUN command rather than by automatic startup. Such a program would issue a CNTACP call to identify itself and then simply wait until a connect request is received from the remote program.

COBOL Call Format:

CALL "CNTACP" USING *stat*, *name*, *objtyp*, *ppn*, *pwd*, *act*

Argument Descriptions:

stat

The *stat* argument is a one-word COMP data item (see Section 3.1). It is set to zero if the call completes successfully. If an error occurs, *stat* is set to one of the error codes listed under "Returned Status" at the end of this section.

name

The *name* argument is a 6-character alphanumeric data item. It identifies the calling program in the network environment and it must be unique to the node. If the remote program addresses this one by object type code and this program is not to be identified by name, this argument can be all spaces.

In any case, the name, if used, consists of 1 to 6 uppercase, alphanumeric characters or the special characters "\$" (dollar sign), "." (period), and "_" (underline). It must be left-justified, with space fill, in the 6-character field. This happens automatically with the VALUE IS clause or the MOVE statement. Leading or embedded spaces are invalid.

objtyp

The object type code (*objtyp*) argument is a 3-character alphanumeric data item expressing an object type code within the range of 0 to 255. If nonzero, it identifies the calling program when the remote program's connect request addressed this program by object type code.

If the remote program addresses this one by name, the *objtyp* argument can be an alphanumeric 0 character. This indicates that object type code addressing is not used. Only one copy of the program can be executed at a time in this case.

The object type code must be left-justified in the data item, with space fill. This happens automatically with the VALUE IS clause or the MOVE statement.

ppn

The *ppn* argument is a 16-character alphanumeric data item. It is a project-programmer number (referred to on some systems as a user identification code, or UIC). The *ppn* is one of three arguments that, if specified, is checked against the corresponding value in the incoming connect request. If all three arguments match, the call succeeds. If the *ppn* value specified in the CNTACP call does not match the *ppn* value in the incoming connect request, the connection is rejected and the CNTACP call terminates with a status code of 5 returned in the *stat* argument.

The project-programmer number must be left-justified in the 16-character field, with space fill. This happens automatically with the VALUE IS clause or the MOVE statement.

pwd

The *pwd* argument is an 8-character alphanumeric data item. It is a password — one of three arguments checked against corresponding values in the incoming connect request. If all three arguments match, the call succeeds. If the *pwd* argument does not match the *pwd* argument in the incoming connect request, the connection is rejected and the CNTACP call terminates with a status code of 5 returned in the *stat* argument.

The password must be left-justified in the 8-character field, with space fill. This happens automatically with the VALUE IS clause or the MOVE statement.

act

The *act* argument is a 16-character alphanumeric data item. It is an account number — one of three arguments checked against corresponding values in the incoming connect request. If all three arguments match, the call succeeds. If the *act* argument does not match the *act* value in the incoming connect request, the connection is rejected and the CNTACP call terminates with a status code of 5 returned in the *stat* argument.

The account number must be left-justified in the 16-character field, with space fill. This happens automatically with the VALUE IS clause or the MOVE statement.

Returned Status:

| <i>stat</i> | Meaning |
|--------------------|---|
| 0 | The CNTACP call completed successfully. |
| 1 | The program already has an existing logical link. Only one link at a time can be used with the concise COBOL interface. The current link must be disconnected before a CNTACP call can succeed. |
| 2 | The format of the call is incorrect. For example, the wrong number of arguments was specified or their lengths are not in the proper range. |

(continued on next page)

| stat | Meaning |
|-------------|--|
| 3 | The CNTACP call was issued by a nonprivileged program. A RSTS/E program must be privileged when it executes a CNTACP call. |
| 5 | The CNTACP call was rejected because the <i>ppn</i> , <i>pwd</i> , or <i>act</i> values in the incoming connect request did not match the values of the corresponding arguments in the CNTACP call. |
| 6 | The physical path to the remote node went down while the connection sequence was in progress. |
| 13 | The alphanumeric <i>name</i> used in the <i>rname</i> argument is already in use on the local system. As noted in Section 2.2, only one program at a time can use a particular name on the local system. |
| 14 | The CNTACP call failed for some reason that will require human intervention to fix. For example, NSP has not been enabled and DECnet is not working at the local node; or the program (job) has been removed as a receiver (denied the use of DECnet) by some privileged user on the system. |

Example:

The following example accepts a connect request. (It would accept the connection requested by the example shown for a CNTCON call in the previous section.)

(Data Division)

```

01      STAT PIC 9(4) COMP.
01      NET-NAME PIC X(6) VALUE IS "INUTRY".
01      OBJTYP PIC X(3) VALUE IS "0".
01      PPN PIC X(16) VALUE IS "1,222".
01      PSWD PIC X(8) VALUE IS "MYNAME".
01      ACCOUNT PIC X(16) VALUE IS "34C".
      .
      .
      .

```

(Procedure Division)

```

      CALL "CNTACP" USING STAT, NET-NAME, OBJTYP, PPN,
      PSWD, ACCOUNT.

```

3.4 Send Record Call (CNTSND)

The CNTSND call sends up to 32767 bytes of data to a remote program over an established logical link. Control returns to the program when the send is complete (that is, all data has been sent and received by the remote program) or when the logical link is disconnected.

Because a CNTSND call allows a large amount of data to be sent, its use requires some care. On a RSTS/E system, for example, network data coming in over a logical link is stored in system buffer space until the program receives the data into its own buffer space. If the remote program never receives the data, and there is not enough system buffer space available to hold the entire message, the remote system will inhibit transmission by the local program over the link until the congestion at the remote system is alleviated. This condition is known as "backpressure." (See Section 6.2.2 for a full discussion of backpressure flow control.) Since backpressure is a dynamic network condition, the subprogram that processes the CNTSND call is designed to continually attempt to send the entire message. If the remote program cannot receive the data being sent, the local CNTSND subprogram will continue trying to send indefinitely.

One reason the remote program might not receive the data is that it is also using the CNTSND call to transmit a large message while the local system has inhibited incoming message flow on its end of the link. Since neither CNTSND call can terminate until the entire message is sent, both programs will wait indefinitely and a deadlock condition will result.

Thus, when two programs communicate using the concise COBOL interface, one should be a sender and the other a receiver, or they should establish some protocol of their own for timing send and receive requests.

COBOL Call Format:

CALL "CNTSND" USING *stat*, *buffer*, *buflen*

Argument Descriptions:

stat

The *stat* argument is a one-word COMP data item (see Section 3.1.1). It is set to zero if the call completes successfully. If an error occurs, *stat* is set to one of the error codes listed under "Returned Status" at the end of this section.

buffer

The *buffer* argument is a data-name referring to the first character position (byte) of a buffer (record) from which the data is to be sent. The data-name can refer to a single data item or an aggregate of data items (for example, a level 01 group name). Any and all PICTURE categories can be used in the fields referenced by data-name.

buflen

This argument is a one-word COMP data item defining the number of bytes in the message. The value of this argument must take into account how the message data in the buffer has been defined. (See the description for the USAGE clause in the *PDP-11 COBOL Language Reference Manual* to see how much computer storage COBOL data items take.) Acceptable values for *buflen* range from 0 through 32767.

Returned Status:

| stat | Meaning |
|-------------|---|
| 0 | The CNTSND call completed successfully. |
| 1 | The program has not established a logical link. A CNTCON or CNTACP call must be executed before the CNTSND call can succeed. |
| 2 | The format of the call is incorrect. For example, the wrong number of arguments was specified. |
| 7 | The logical link was disconnected before the complete message was sent. This could occur because the remote program disconnected the link or because of some network or system malfunction (the physical path went down, the remote system crashed, and so forth). |
| 14 | The CNTSND call failed for some reason that will require human intervention to fix. For example, NSP has not been disabled and DECnet is not working at the local node; or the program (job) has been removed as a receiver (denied the use of DECnet) by some privileged user on the system. |

Example:

(Data Division)

```
01    PAYREC.
      02    EACH-ONE OCCURS 20 TIMES.
          03    NAME.
              04    LAST-NAME PIC X(20).
              04    FIRST-NAME PIC X(20).
              04    MID-INIT PIC X(1).
          03    HOURS PIC X(3).
          03    HOURLY-WAGE PIC X(4).
01    STAT PIC 9(4).
01    BUFLN PIC 9(4) COMP VALUE IS 960.
      .
      .
      .
```

(Procedure Division)

```
CALL "CNTSND" USING STAT, PAYREC, BUFLN.
```


3.5 Receive Record Call (CNTRCV)

The CNTRCV call is used to receive a message from a remote program. The call defines a buffer area where the message is to be stored, the size of the buffer in bytes, and a COMP data item that will be set to the number of bytes in the received message. Control returns to the program when the receive is complete or the logical link is disconnected.

Note that it is possible to suspend program execution indefinitely on a CNTRCV call if the remote program never sends a message to be received and the logical link is never disconnected. As discussed with the CNTSND call, two programs using the concise COBOL interface should be designed such that one is always a sender and the other a receiver, or they should establish their own protocol for send and receive requests.

COBOL Call Format:

CALL "CNTRCV" USING *stat*, *buffer*, *buflen*, *datlen*

Argument Descriptions:

stat

The *stat* argument is a one-word COMP data item (see Section 3.1.1). It is set to zero if the call completes successfully. If an error occurs, *stat* is set to one of the error codes listed under "Returned Status" at the end of this section.

buffer

This argument is a data-name referring to the first character position (byte) of a buffer to contain the received message. The data-name can refer to a single data item or an aggregate of data items (for example, a level 01 group name). Any and all PICTURE categories can be used in fields referenced by data-name. The format should, of course, reflect the type of data sent from the remote program.

buflen

The *buflen* argument is a one-word COMP data item defining the length of the buffer in bytes. Acceptable values range from 0 through 32767. If the received message contains more bytes than indicated by the *buflen* argument, the remainder of the message will be discarded.

datlen

The *datlen* argument is a one-word COMP data item. When the call completes, it will contain the actual number of bytes transferred to the area defined by the *buffer* argument.

Returned Status:

| stat | Meaning |
|------|--|
| 0 | The CNTRCV call completed successfully. |
| 1 | The program has not established a logical link. A CNTCON or CNTACP call must be executed before a CNTRCV call can succeed. |
| 2 | The format of the call is incorrect. For example, the wrong number of arguments were specified in the call. |
| 7 | The logical link was disconnected before the entire message was received by the local NSP. Some DECnet user interfaces (including the DECnet/E extended COBOL interface) allow a program to abort a logical link — that is, to disconnect the logical link regardless of whether an entire message has been transferred. This may have happened; or the link may have been disconnected because the line went down or the remote system crashed. |
| 8 | The logical link no longer exists because the remote program has disconnected the link. This is regarded as an end-of-file for concise COBOL programs. |
| 9 | The length of the message received exceeded the number of bytes defined by the <i>buflen</i> argument. The remaining data was discarded. |
| 14 | The CNTRCV call failed for some reason that will take human intervention to fix. For example, NSP has disabled and DECnet is not working at the local node; or the program (job) has been removed as a receiver (denied the use of DECnet) by some privileged user on the system. |

Example:

The following example receives up to 1000 bytes of data from the network over an established logical link, storing it in a data record named INBUF.

(Data Division)

```
01      INBUF PIC X(1000).  
01      STAT PIC 9(4) COMP.  
01      BUFLN PIC 9(4) COMP VALUE IS 1000.  
01      DATLEN PIC 9(4) COMP.
```

·
·
·

(Procedure Division)

```
CALL "CNTRCV" USING STAT, INBUF, BUFLN, DATLEN.
```

3.6 Disconnect Call (CNTDIS)

The CNTDIS call disconnects a logical link. If the remote program uses the DECnet/E concise COBOL interface and is executing a CNTRCV call, that CNTRCV call will terminate with the end-of-file status value (*stat* = 8).

COBOL Call Format:

CALL "CNTDIS" USING *stat*

Argument Description:

stat

The *stat* argument must be a one-word COMP data item. It is set to zero if the call completes successfully. If an error occurs, *stat* is set to one of the error codes listed under "Returned Status."

Returned Status:

| <i>stat</i> | Meaning |
|-------------|--|
| 0 | The CNTDIS call completed successfully. |
| 1 | The program has not established a logical link to be disconnected. A CNTCON or CNTACP call must be issued before a CNTDIS call can succeed. |
| 2 | The format of the call is incorrect. For example, the wrong number of arguments was specified. |
| 10 | There are still messages waiting to be received with a CNTRCV call. A program cannot disconnect the link until it has processed all the messages that have been sent to it. |
| 11 | The link has already been disconnected. |
| 14 | The CNTDIS call failed for some reason that will take human intervention to fix. For example, NSP has been disabled and DECnet is not working at the local node; or the program (job) has been removed as a receiver (denied the use of DECnet) by some privileged user on the system. |

Example:

The following example disconnects the program's logical link.

(Data Division)

```
01      STAT PIC 9(4) COMP.  
          .  
          .  
          .
```

(Procedure Division)

```
CALL "CNTDIS" USING STAT.
```


Part III
**Network Programming with the
Extended COBOL Interface**

Chapter 4

Extended Discussion of Logical Links

This chapter expands the discussion of the logical link — the basic element of DECnet task-to-task communication. As mentioned in Chapter 2, a logical link is a software path for the exchange of data between two programs running in a network.

Throughout the remainder of this chapter, and the two chapters that follow, various logical link parameters are discussed — object name and type code, local and remote link address, flow control data request counts, and so on. These parameters can be observed on an operational node by issuing the NCP SHOW LINK command (for a single link) or the SHOW ACTIVE LINKS command (for all current links). This can be very helpful when debugging a network program. See the *DECnet/E System Manager's Guide* for details on using these NCP commands.)

4.1 Creating a Logical Link in the Extended COBOL Interface

Creating a logical link is a cooperative venture. Two programs must agree to communicate before a logical link is established. The procedure is basically the same for all implementations of DECnet. The program that initiates the request for a logical link connection is called the source. The other program is called the target. This distinction is made only during the connection sequence.

NOTE

Other network implementations make the source program the “link master.” The source must poll the target for data and only the source can disconnect the link. DECnet does not make this distinction. Once the connection is established, the terms “source” and “target” have no significance and both the local and remote programs have equal access to the logical link.

The request for a connection takes the form of a Connect Initiate Message that includes detailed information identifying the remote node and target program to which the connection is to be made. It also establishes an identifier for the link itself. In DECnet/E, the link identifier is called a user link address, or ULA. A ULA is a number from 1 to 255 that is used by the source program to refer to the link during subsequent send and receive operations.

The NSP software at the source node does some initial checking of the request (to determine if it recognizes the remote node, for example). If all is well, the source NSP forwards the connect request to the NSP at the target node. The target NSP does its own checking to determine if the connection can be made. If the connection cannot be made, the target NSP rejects the connection. Otherwise, it delivers the connect request to the target program. The target program can then choose to accept or reject the logical link.

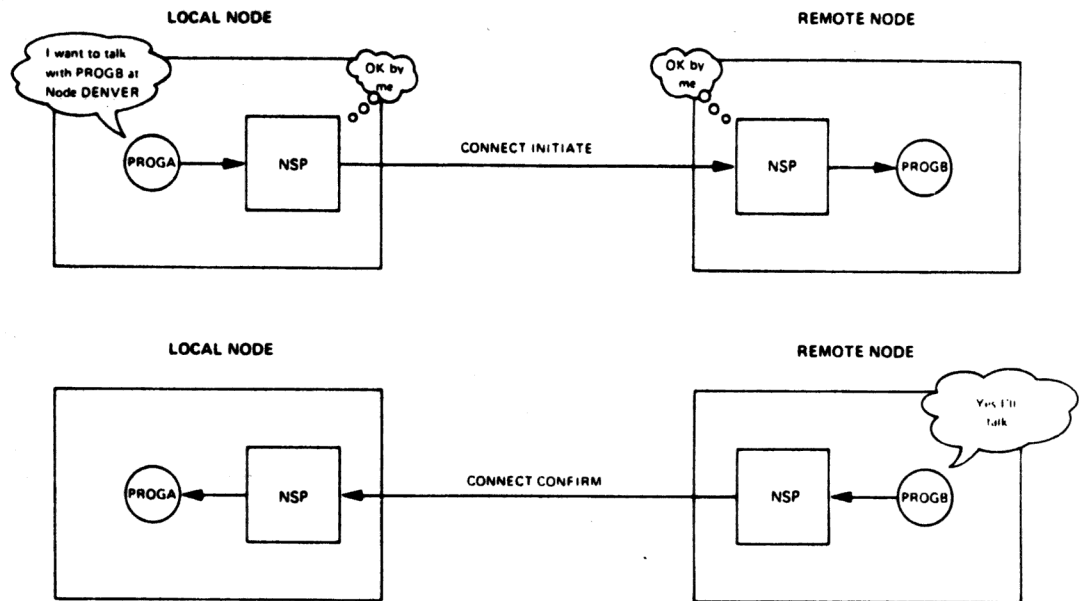
If the target program accepts the link, it selects its own link identifier by which it will refer to the link. The two link identifiers are unrelated to one another. They can even take quite different forms if the programs are written in different languages or for different operating systems. If the remote node is a DECnet/E node, this identifier is again called a user link address and can range from 1 to 255. The number selected is completely independent of the selection made by the source program. The source and target NSPs relate the two identifiers to the same logical connection and keep the data exchanged over the link separate from data exchanged over other logical links.

The target program's confirmation or rejection of the link is passed back over the network to the source NSP, which passes this information on to the source program. This confirmation or rejection takes the form of a Connect Confirm or Connect Reject Message that is queued for the source to receive and process.

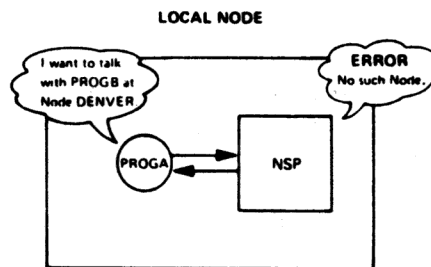
If the connection is confirmed after this mutual handshaking, the two programs can then send and receive data over the link, each referring to the link by the identifier it assigned during the connection sequence. Both programs can send data simultaneously, since the DECnet software provides logical full-duplex transmission regardless of the characteristics of the intervening network. Either program can break the logical link connection when it is no longer needed.

In Figure 4-1, A shows the procedure for establishing a logical link connection. The target NSP considers the link "up" (that is, it will allow the target program to send data) when the source NSP has acknowledged receipt of the Connect Confirm Message. The source NSP considers the link up when it receives the Connect Confirm Message. In Figures 4-1, B, C, and D illustrate the three ways in which a request for a link might be rejected.

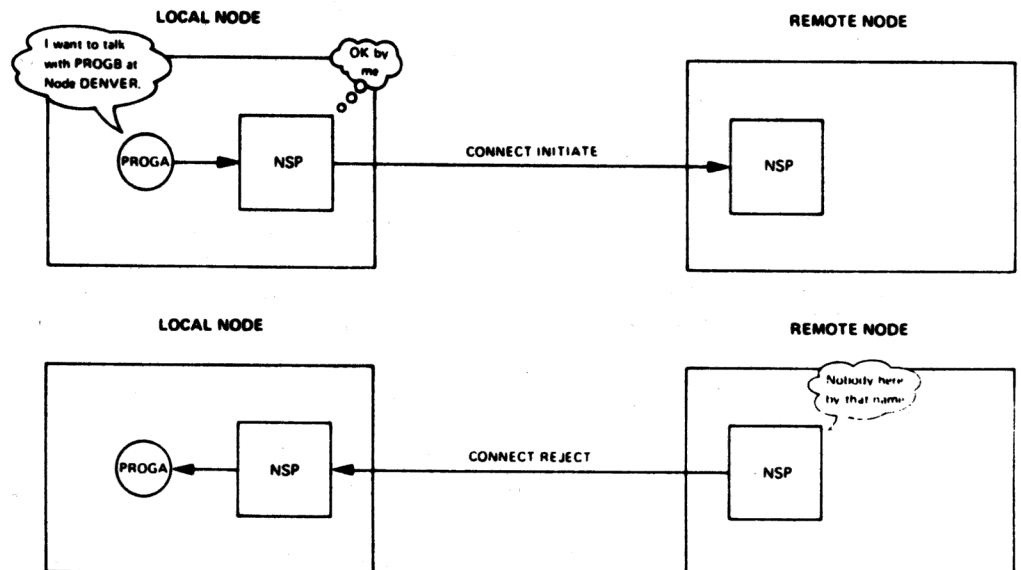
A. Remote program confirms logical link connection



B. Local NSP returns an error on connect request



C. Remote NSP rejects logical link connection



(continued on next page)

Figure 4-1: Logical Link Connections

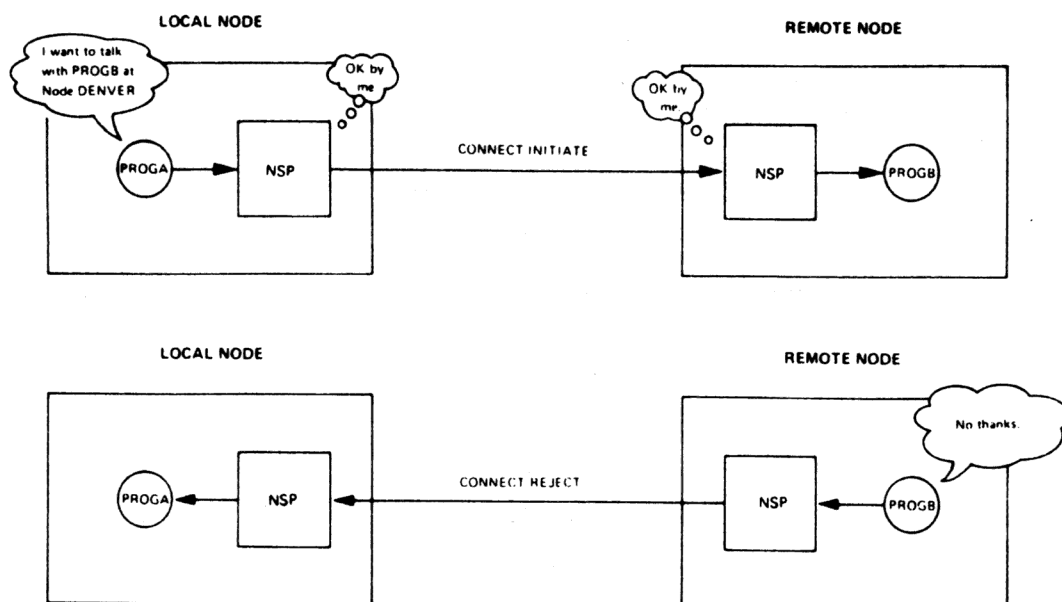


Figure 4-1 (Cont.): Logical Link Connections

4.2 Advantages of Logical Links

In addition to making it possible to sort out the many physical streams of data that pass over a single physical line, the logical link structure maintained by NSP has several other advantages.

- Reduction of addressing data required for message transmission, thereby providing efficient line usage
- Separation of data into separate message streams
- Provision for high-priority "Interrupt" messages
- Ability to control the flow of data, thereby relieving possible network congestion

4.2.1 Efficient Line Usage


Once established, a logical link simplifies the addressing necessary for the exchange of data between two programs, both at the user's level and in transmission of data across the network. As mentioned previously, a link identifier is assigned by each program to identify an established logical link. This identifier is more convenient than the longer, detailed address used in identifying the remote node and program in the original Connect Initiate Message.

During a connection sequence, the source NSP and target NSP also establish a set of addresses that they use to identify the logical link. These addresses are called the local link address (LLA) and the remote link address (RLA), respectively. Each NSP refers to its own as the local link address and the other as the remote link address.

While individual ULAs must be unique to the source program and to the target program, they are by no means required to be unique throughout the network, or even within the local node. Thus, the LLA and the RLA are required by NSP to separate the many distinct message streams that are multiplexed onto the communications lines. Since each NSP's LLA is unique to its own node, these NSP addresses reduce the amount of control and identifying information that must accompany the user data transferred across the link (see Figure 4-2).

Figure 4-3 illustrates how these addresses work, showing two DECnet/E nodes with two established logical links. Program A, for example, has established a logical link with Program D. Program A refers to the link with a user link address of 1. Program D refers to the link with a user link address of 23. Each NSP has established its own local link address. For Program A's NSP, the LLA is 123456. For Program D's NSP, the LLA is 024602. Each NSP keeps track of the other's LLA as a remote link address.

A. Message sent across the line for a connect request



B. Data message sent across the line over a logical link




Figure 4-2: Logical Links Provide Line Efficiency

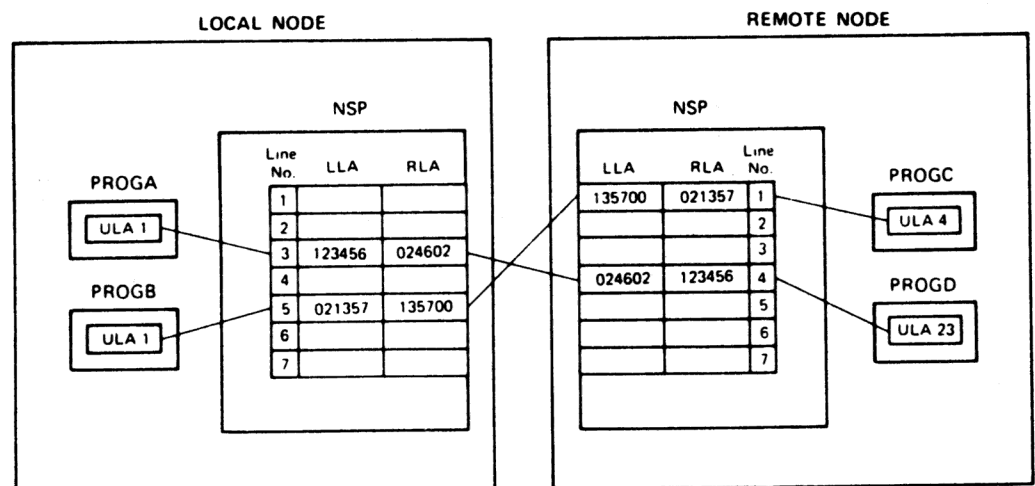


Figure 4-3: User Link Addresses, Remote Link Addresses, and Local Link Addresses for Logical Links

4.2.2 Separate Data by Intended Use

Logical links provide a means by which incoming messages can be separated into different message streams. A program can establish different logical links to communicate with different remote programs. Or it can establish several logical links with the same remote program to exchange data intended for different purposes.

Before a program can use the message services, it must first register this intention with NSP. This involves declaring the name by which the program is to be known to other remote programs in the network (that is, its identity), as well as any limits or restrictions that will affect the way in which the program can use the message services.

When a program declares its intent to use the network message services, NSP sets up a receive queue to hold pending messages. Thereafter, all messages received for the program are placed in the queue. Messages from different logical links are interspersed in the queue but the receiving program can retrieve either the first message in the queue or the first message from a particular logical link.

The concept of separate message streams for separate logical links is illustrated in Figure 4-4.




Figure 4-4: Data Is Queued for Logical Links in Separate Streams

4.2.3 Interrupt Messages

Special high-priority Interrupt Messages can be sent over logical links in all DECnet implementations. DECnet/E queues incoming Interrupt Messages at the head of the pending message queue, behind the first message in the queue (since the user program could be in the process of retrieving the first message) and behind any other pending Interrupt Messages queued for the program.

In the DECnet/E implementation, these Interrupt Messages are not true interrupts in the sense of causing an immediate jump to an interrupt processing routine. (Some DECnet implementations *do* treat Interrupt Messages as true interrupts.) They are simply messages that can be identified as different from ordinary data messages and that are placed at the head of the pending message queue (see Figure 4-5).

If a program is designed to send Interrupt Messages, the receiver at the other end of the link must be set up to process them. Interrupt Messages can contain a small amount of user data so that programs can be designed to recognize Interrupt Messages for different purposes.

Interrupt Messages are discussed in more detail in Chapters 5 and 6.




Figure 4-5: Queuing of Interrupt Messages for Logical Links

4.2.4 Flow Control

NSP maintains each program's queue of pending messages in system buffer space. A message is stored in the queue until it is received by the user program. Thus, pending messages tie up valuable system resources. Several flow

control mechanisms are used in DECnet to protect systems from being flooded with messages. Some of these mechanisms are system-regulated while one mechanism is selected and controlled by the user program.

Each receiver program can select one of the following three types of program-regulated flow control:

- Segment flow control
- Message flow control
- No active flow control

A discussion of program-regulated flow control is given in Section 6.2.3.

4.3 DECnet/E Logical Link Limitations

The implementation of DECnet for RSTS/E allows a maximum of 127 active logical links for the entire node at any one time. The system manager can reduce this maximum to some number between 1 and 127. There is no limit to the number of logical links per program other than the system maximum.

Chapter 5

Overview of Extended DECnet/E Message Services

RSTS/E provides message services for communication between local programs. The DECnet/E message services are implemented as an extension of these local services. In addition, DECnet/E offers new capabilities for local communication not found with the original RSTS/E message mechanism. The original local communication services are described in the *RSTS/E Programming Manual*. They are repeated in this document both for completeness and because they are the logical basis for the DECnet extensions.

We emphasize here that it is the way in which the user interfaces with DECnet from RSTS/E that is implemented as an extension of local capabilities. An RSX-11 programmer might interface with DECnet somewhat differently. The underlying architecture provides the consistency. The programmer at each system uses a familiar language and form. DECnet handles both the transmission itself and the translation necessary at both nodes to convert the original call to a form that can be transmitted over the line and back to the format expected at the receiving node.

Table 5-1 lists the calls used for interprogram communication by programs running under DECnet/E. The remaining sections of this chapter give a general description of each of the calls. Specific COBOL calls for the extended interface, with detailed formats, are given in Chapter 7.

Table 5-1: Summary of Calls for Interprogram Communication

| Call | Function |
|---------------------------|---|
| Declare Receiver | Registers the program with the operating system for message services. |
| Remove Receiver | Terminates message operations for the program. |
| Get Local Node Parameters | Obtains local node information. |
| Log User Event | Queues a user-generated event to the local event processor for logging. |
| Send Local Data Message | Transmits user data to a local program. |

(continued on next page)

Table 5-1 (Cont.): Summary of Calls for Interprogram Communication

| Call | Function |
|-------------------------------|---|
| Send Connect Initiate Message | Requests a logical link connection with another program. |
| Send Connect Confirm Message | Accepts a logical link connection requested by another program. |
| Send Connect Reject Message | Rejects a logical link connection requested by another program. |
| Send Network Data Message | Transmits user data to a network program over an established logical link. |
| Send Interrupt Message | Transmits interrupt data to a network program over an established logical link. |
| Send Link Service Message | Requests data over a flow controlled logical link. |
| Send Disconnect Message | Disconnects an established logical link, after all pending messages have been sent. |
| Send Link Abort Message | Disconnects an established logical link immediately, destroying any messages waiting to be sent. |
| Receive | Receives a message from the queue of pending messages. (One of nine message types, corresponding to the nine Send calls listed above, could be received.) |

5.1 General Remarks

The Declare Receiver call informs the local RSTS/E monitor that the program issuing the call intends to communicate with other programs. This call identifies the program and defines which of the services will be required. The Remove Receiver call is issued when message operations are complete for the program and releases system resources used for message services. The Get Local Node Parameters call returns information to the calling program concerning the local node. The Log User Event call permits a user-written program to queue an event to the event processor for logging.

Beyond these, the operations available fall into two categories: send and receive. The entities sent and received are called messages. They consist of control information identifying the type of message, and in most cases, data of interest to the programs exchanging information. Nine message types are available for these operations.

A separate call is used to send each of the message types. That is, there is a Send Connect Initiate Message call, a Send Local Data Message call, and so forth. When one of these calls is executed, NSP uses the parameters and user data included in the call to format the appropriate network message into system buffer space. Necessary header information is also included. This message is then queued to Transport for transmission across the network. The way in which a message is queued for transmission depends on the type of message. (A discussion of transmission queuing is given in Section 6.2.1.)

The Receive call retrieves a message from the program's queue of pending messages maintained by NSP in system buffer space. NSP queues incoming messages for the program as they come in across the network. Messages from local senders are also placed in the pending message queue.

NOTE

Within the current context, *local* programs are programs running at the local node that send messages using only the Send Local Data Message call. *Network* programs, on the other hand, are programs running either locally or remotely that use DECnet to establish and send data over logical links.

The Receive call returns control information identifying the sender and the type of message (one of the nine listed previously in Table 5-1). User data accompanying the message is returned to a user buffer specified in the Receive call where it can then be examined and processed.

A user program can issue general or selective Receive calls. On a general Receive, the first message in the queue is returned. A selective Receive can select the first message from any local sender, a particular local sender, any network sender, or a particular logical link.

There can be times when no pending messages are in the queue when a Receive call is issued. One option available in the Receive call causes the program to "sleep", or suspend execution, until some action takes place — until a message is queued or a specified amount of time has passed, for example. The Receive call can also be issued to retrieve only part of a message. A program with limited buffer space can choose to retrieve only part of the user data with one Receive call. Subsequent Receive calls will retrieve the rest of the data, or the remaining data can be discarded.

5.2 Registering with the Monitor: Declare Receiver Call

A program declares its intent to use the message services by issuing a Declare Receiver call. This call must be executed before a program can send network messages or receive messages of any type. A program can, however, send local data messages without executing this call.

The Declare Receiver call does more than just declare intent, however. It also defines the way in which other programs throughout the network are to address this program (the program's identity) and the way in which the program can thereafter use the message services.

The Declare Receiver call has some rather far-reaching effects. For this reason, the parameters specified with the call and how the system uses them are described here in detail.

5.2.1 Declaring Identity

A Declare Receiver call identifies the calling program with an object type code, a name, or both. The manner in which a program declares its identity affects the way incoming link connections are handled.

A name is an ASCII identifier that can be used by either local or network programs to refer to the declarer. No two programs at the local node can use the same name at the same time. If a name is given, it must be unique within the local node.

An object type code, on the other hand, is simply a number in the range of 1 to 255. It can be used by other network programs to refer to the declarer. (Local programs cannot address a receiver by object code.) These codes simplify the identification of programs performing identical functions at different nodes of a network, without going through the cumbersome process of arriving at and enforcing standard naming conventions.

An object code also need not be unique, and more than one receiver can declare identity with the same object code simultaneously. A program performing a general-purpose function of interest to the entire network can incur heavy use. By permitting more than one program to declare identity with the same object code at the same time, DECnet/E permits execution of multiple copies of a program to meet sporadic heavy usage.

A range of object type codes (1 through 127) is reserved for DECnet use. The remainder (128 through 255) are available for user applications programs. Currently assigned object type codes for DECnet use are given in Appendix A.

NOTE

In the DECnet environment it is recommended that, whenever possible, programs declare their identity with and be addressed by nonzero object type codes.

Network addressing is discussed in detail in Chapter 6.

5.2.2 Declaring Intended Usage

In addition to identifying the program, the Declare Receiver call defines how the program intends to use the message services. For example, a program can limit the incoming messages to those from local senders only. Or, it can limit incoming messages to those from network senders only.

Other protective limits can be set in the Declare Receiver call.

- The maximum amount of system buffer space to be used for the pending message queue
- The maximum number of messages that can be stored in the pending message queue
- The maximum number of active logical links that the program can have at any given time

These limits affect incoming messages and incoming requests for logical links. The Declare Receiver call does not limit the type or number of send calls a program can execute.

5.2.3 DECnet/E Use of the Declare Receiver Call

When a Declare Receiver call is issued, the system sets up a 32-byte Receiver ID Block (RIB) to associate the information with the RSTS/E job number. The Receiver ID Block actually applies to the job rather than the program that issues the Declare Receiver call. This distinction is rather fine for most applications, but does make it possible, for example, to chain together a set of programs using message services. The first program in the series would execute a Declare Receiver call. Then other programs in the chain would also be valid receivers as long as they were always executed under the same job number.

The Receiver ID Block holds the arguments passed with the Declare Receiver call, along with other system- and NSP-maintained information, such as the address pointers of the pending message queue and the address of the list of active logical links.

5.3 Remove Receiver Call

A Declare Receiver call remains in effect until a Remove Receiver call is issued. The Remove Receiver call releases the Receiver ID Block. All pending messages are discarded and all active logical links are destroyed. (NSP sends a Link Abort Message to the remote programs for any active logical links.)

Like the Declare Receiver call, the Remove Receiver call applies to the RSTS/E job. If a program does not issue a Remove Receiver call before terminating execution, other programs executed later under this job number will be unable to issue a Declare Receiver call. (As far as the system is concerned, the old one is still in effect.) A Remove Receiver call should be issued as soon as message services are no longer needed. This practice prevents queuing of unwanted messages and releases the system buffer space used for the Receiver ID Block.

Also, any program that could be stopped by a terminal user's **CTRLC** should have provisions to issue a Remove Receiver call and terminate gracefully. If a job is killed, NSP automatically destroys any active logical links, discards messages, and releases any system resources used for message operations.

NOTE

It is generally good practice for a program to issue a Remove Receiver call immediately before issuing a Declare Receiver call. This prevents errors on the Declare Receiver call caused by abnormal termination of a previous program running under the same job number.

5.4 Get Local Node Parameters Call

The Get Local Node Parameters call returns information to the calling program concerning the local node. This information consists of the node name, address, and alias (if any), and the node's default project-programmer number (if any). (See the *DECnet/E System Manager's Guide* for a discussion of alias.)

A general-purpose network program — one that runs on several different nodes rather than being specific to any one node — would use this call to obtain the node name, number, and alias for use in the connection sequence. The program uses this information to identify itself to the remote node and program. A program that is started automatically as the result of an incoming connect request would use this call to obtain the default project-programmer number for use in validating the connection and logging into a user account. (See “Automatic Job Startup” in Chapter 6 for a discussion of the default account.)

5.5 Log User Event Call

The Log User Event system call permits a user-written program to queue an event to the system event processor for logging. Before the event is logged, it is time-stamped by the system in the standard Network Information and Control Exchange (NICE) protocol.

Optional parameter data can be passed to the event processor for processing and output. If it is included, this data must be in NICE protocol format. (See the *DIGITAL Network Architecture, Network Management Functional Specification*, Version 3.0.0, for further information on event logging.)

NOTE

This system call performs a highly specialized function which requires a great deal of special knowledge of DECnet and the DIGITAL Network Architecture. It is provided as a convenience for the sophisticated network user and is not intended for normal network programs.

5.6 Send/Receive Local Data Message

The Send Local Data Message call transfers up to 532 bytes of user message data to a local program. The message is placed in the appropriate receive queue of the target program. The target program can be identified either by job number or by the name given in its Declare Receiver call.

A Receive call that returns a Local Data Message returns control information identifying it as a Local Data Message from a specific local sender, and up to 532 bytes of user data.

The Local Data Message is not used for DECnet communication but is included here for reference. A user program can carry on local conversations and maintain one or more DECnet logical links at the same time.

5.7 Send/Receive Connect Initiate Message

The Send Connect Initiate Message call is used to request a logical link connection with a remote program.

The target node and program are identified by a 120-byte Connect Data Block passed as part of the call. The following information is specified with the call:

The user link address (ULA) that the source program will use to refer to this link.

The type of receiver flow control desired by this program — that is, the type of flow control to be imposed on incoming data messages. (A full discussion of flow control is deferred to Chapter 6.)

A receive maximum. The maximum amount of data to be received by the program over this link in one Network Data Message. The size of the receive buffers allocated by NSP will also limit the amount of data that can be passed in a single Network Data Message. If the maximum buffer size set by Network Management for NSP is not large enough to handle the receive maximum as specified by the user program, NSP will alter the receive maximum to the smaller limit before forwarding the Connect Initiate Message. The local program is informed of this change (if any is made) by a field in the Connect Confirm Message from the remote program when it accepts the link.

A Connect Data Block identifying the target node and program to which the connection is requested. The target program can be identified by name or object type code. NSP inserts the local program identification, as specified in the source program's Declare Receiver call, into the Connect Data Block. This allows the remote program to determine who sent the connect request and decide whether to accept or reject the connection. The Connect Data Block can also contain data required by the remote program or system for access control. Such data might consist of a project-programmer number, password, and account number — or whatever is required by the remote program or system.

Up to 16 bytes of user data.

A Receive call that returns a Connect Initiate Message from the queue of pending messages returns control information identifying it as a Connect Initiate Message from a specific network program, and requirements imposed by the remote program. The information includes the following:

The local link address (LLA) assigned by the local NSP to refer to this logical link. The program will use the LLA to refer to the link in its responding Connect Confirm or Connect Reject Message. (Note that the

LLA is only used to refer to a link while no user link address (ULA) is in effect. The ULA is established by the receiving program in its Connect Confirm Message that accepts the connect request and is used thereafter by the program to refer to the link.)

The type of receive flow control desired by the remote program — that is, the type of flow control that will be imposed on outgoing data messages over this link.

A receive maximum. The maximum amount of user data that can be received by the local program over the link. This maximum is calculated by the local NSP, based on the size of receive buffers it allocates. The program can reduce this maximum further in its responding Connect Confirm Message.

A transmit maximum. The maximum amount of user data that can be transmitted by the program over this logical link in one Network Data Message. This maximum is imposed by either the remote NSP or the remote program and corresponds to the remote program's receive maximum.

A 120-byte Connect Data Block, identifying the remote node and program requesting the link. The Connect Data Block also indicates how the remote program addressed this one (by name or object type code). If the program requires access control information from the remote program, such information would also be included as part of the Connect Data Block. The Connect Data Block and up to 16 bytes of user data are delivered to the user buffer specified in the Receive call.

5.8 Send/Receive Connect Confirm Message

The Send Connect Confirm Message call is used to accept a logical link requested by a remote program. That is, it is a positive response to a received Connect Initiate Message. The following information is specified with the call:

The user link address (ULA) by which the program will hereafter refer to this link.

The local link address (LLA) that identifies the link being accepted. (The LLA is taken from the Connect Initiate Message received for this particular link.)

The type of receive flow control desired by the calling program for incoming data messages over this link. (Flow control options are discussed in detail in Chapter 6.)

A receive maximum. The maximum amount of user data that the program wishes to receive as a unit (that is, as one Network Data Message).

Up to 16 bytes of user data.

A Receive call that returns a Connect Confirm Message from the pending message queue returns control information identifying it as such, along with requirements for messages to be transmitted over the link. The information includes the following:

The user link address (ULA) established by the receiving program in the Connect Initiate Message that requested the link. This identifies the particular link being confirmed.

The type of receive flow control requested by the remote program — that is, the type of flow control imposed on outgoing data messages. (Flow control options are discussed in detail in Chapter 6.)

A receive maximum. The maximum amount of user data that the local program will receive on this link in one Network Data Message. This will be the same or less than the size requested in the program's Connect Initiate Message. A smaller size is indicated by the local NSP if it cannot handle the size requested in the Connect Initiate Message.

A transmit maximum. The maximum amount of data the local program can send to the remote program as a unit (with one Network Data Message) over this link.

Up to 16 bytes of user data, delivered to the user buffer specified in the Receive call.

5.9 Send/Receive Connect Reject Message

A Send Connect Reject Message call is used to reject a logical link requested by a remote program. That is, it is a negative response to a received Connect Initiate Message. The following information is specified with the call:

The local link address (LLA) assigned by the local NSP to identify the link being rejected. (The local link address is taken from the received Connect Initiate Message.)

A code identifying the reason for the rejection. Ordinarily, this code is zero unless the program is rejecting the connection due to some standard reason, such as "invalid accounting data."

Up to 16 bytes of user data.

A Receive call that returns a Connect Reject Message from the queue returns control information identifying it as such and the following information:

A user link address identifying the link being rejected. This user link address will be the same as that established by the local program in the Connect Initiate Message that requested the link. The link could have been rejected by either the remote NSP or the remote program.

A code identifying the reason for the rejection if the link was rejected by the remote NSP or by the remote program for some standard reason, such as "invalid accounting data." (A list of NSP rejection codes is found in Appendix B.)

Up to 16 bytes of user data, if the link was rejected by the remote program for some nonstandard reason. This data will be delivered to the user buffer specified in the Receive call.

5.10 Send/Receive Network Data Message

The Send Network Data Message call transmits user information to a remote program over an established logical link. The logical link is identified by the user link address (ULA) established by the calling program in its Connect Initiate or Connect Confirm Message. The amount of user data that can be sent is limited to the transmit maximum established by the remote program (see the previous discussion of a received Connect Initiate or Connect Confirm Message). Sending a Network Data Message is subject to flow control imposed by the system and by the remote program, as indicated by the flow control option selected by the remote program. A program can request that link status information be returned by a Send Network Data call and use the information returned to schedule subsequent send and receive requests. A full discussion of flow control is deferred to Chapter 6.

A Receive call that returns a Network Data Message returns identifying control information and delivers the user data to a buffer specified in the call. The amount of user data received in one Network Data Message will not exceed the receive maximum established by this program for the logical link (see the previous discussion of a sent Connect Initiate or Connect Confirm Message). Incoming Network Data Messages are subject to system flow control and to the flow control option, if any, requested by the local program in its Connect Initiate or Connect Confirm Message. If the program requests flow control, it must send Link Service Messages to request data from the remote program (see Chapter 6).

5.11 Send/Receive Interrupt Message

The Send Interrupt Message call is used to transmit high-priority data to a remote program over an established logical link. All DECnet systems maintain separate data streams for Interrupt and Data Messages. The Interrupt Message will be delivered to the remote program ahead of all Data Messages waiting to be processed at the remote end. Up to 16 bytes of user data can be sent in an Interrupt Message. Sending an Interrupt Message is subject to flow control imposed by the remote system or program. The local program can request that link status information be returned following a Send Interrupt Message call and use the information returned to schedule subsequent send and receive requests. A full discussion of flow control is given in Chapter 6.

A Receive call that returns an Interrupt Message returns identifying control information and delivers up to 16 bytes of user data to a buffer specified in the call. The Interrupt Message will have been queued at the head of the pending

message queue, behind the first Data Message and behind any other pending Interrupt Messages queued for the program. Incoming Interrupt Messages are subject to flow control imposed by the local NSP. After receiving each Interrupt Message, the local program must send a Link Service Message to reen able incoming Interrupt Messages from the remote program over a logical link (Chapter 6).

5.12 Send/Receive Link Service Message

The Send Link Service Message call is used to request data over a logical link — either to request data from a remote program over a flow controlled link, or to reen able incoming Interrupt Messages from the remote program. In these two cases, detailed link status information can be returned or the call can be issued solely to obtain this information.

A Link Service Message is returned on a Receive call when a link becomes unblocked. If a local program has tried to send a Network Data or Interrupt Message and failed because the remote system or program has inhibited flow, the local NSP informs the local program when the condition clears. To encourage efficient processing of pending messages, NSP returns a Link Service Message on a Receive call only when the pending message queue is empty. Detailed link status information is returned that the local program can use in regulating its send and receive requests. (The format of the status information for a received Link Service Message is the same as that returned, upon request, for a Send Network Data, Interrupt, or Link Service call.)

For a more complete understanding of the uses of Link Service Messages, see the discussion of flow control in Chapter 6.

5.13 Send/Receive Disconnect Message

The Send Disconnect Message call is used to break an established logical link provided that all Network Data, Interrupt, and Link Service Messages previously sent over the link have been acknowledged by the remote NSP. If they have not been acknowledged, the call will terminate with an error and the link will not be broken. Successful completion of the call means that no more messages can be sent over the link. The user link address is freed and can be used to establish another logical link. However, messages already in the pending message queue will remain. They can (and should) be cleared from the queue by doing selective Receive calls using the user link address of the disconnected link. Once the link has been disconnected, new messages received for the link will be discarded by NSP. The Disconnect Message is sent to the remote system. When it is received by the remote NSP, the link is broken at that end. Up to 16 bytes of user data can be sent with a Disconnect Message.

A Receive call that returns a Disconnect Message from the pending message queue returns control information identifying it as such, as well as a user link address and up to 16 bytes of user data. Receipt of a Disconnect Message indicates that the remote program has terminated the logical link. Furthermore, it indicates that all previous Network Data Messages, Interrupt Mes-

sages, and Link Service Messages sent by the remote program were placed in the pending message queue by the local NSP before the Disconnect Message was received. Once the Disconnect Message is received by the local NSP, the local program cannot send any more messages over the link.

The Disconnect Message is useful in "master-slave" communication when the master program only transmits data and the slave program only receives data. Since there is no two-way communication, the master program can issue a Disconnect Message to break the link with the assurance that the remote system has received all the data sent and queued it for the slave program. Successfully completing the disconnect does not guarantee that the slave has processed the data, however. Where such guarantees are desirable, or where two-way communication is desired, programs should agree to terminate the logical link by exchanging Network Data Messages, Interrupt Messages, or whatever.

5.14 Send/Receive Link Abort Message

A Send Link Abort Message is used to break a logical link immediately. Unlike a disconnect, no attempt is made to ensure that messages previously sent over the link have been acknowledged by the remote system. In fact, NSP discards any messages it has queued for transmission for the link. (Section 6.2.1 describes NSP transmit queuing.) No further data can be sent over this logical link. Received messages already in the pending message queue will remain but new messages received after the Link Abort Message is issued will not be queued for the link. Messages already in the queue can (and should) be cleared by doing selective Receive calls using the user link address of the aborted link.

If the link is established, the Link Abort Message is sent to the remote system as notification that the link is broken and up to 16 bytes of user data can accompany the Link Abort Message. A Link Abort Message can also be sent for a logical link that has not yet been confirmed by the remote program. In this case, NSP breaks the link at the local node, but it does not notify the remote system. (Hence, the user data does not reach the remote program.) Should the remote program confirm the link, the local NSP no longer knows the link and it so informs the remote NSP, which notifies the remote program. Should the remote program reject the link, the local NSP informs the remote NSP that it has no such link but since the remote program rejected the link anyway, no notification is given to the remote program.

A Receive call that returns a Link Abort Message from the pending message queue returns control information identifying it as such, as well as a user link address and a local link address to identify the link aborted. The local link address identifies a link aborted during the initial connection sequence. For example, if the local NSP received a Connect Initiate Message, but the local program has not yet responded with a Connect Confirm or Connect Reject Message, a Link Abort Message is queued with the same local link address as was specified in the received Connect Initiate Message. The link could have been aborted by the remote NSP or the remote program. If the link was aborted by the remote NSP, the message will contain a reason code identify-

ing the reason for the rejection. If the link was aborted by the remote program, up to 16 bytes of user data can also be delivered with the Link Abort Message. Once the local NSP receives the Link Abort Message, the local program cannot send any more messages over the link.

Chapter 6

Network Addressing and Flow Control

With the background given in Chapters 4 and 5, we can now discuss two areas of interprogram communication in a DECnet network: network addressing and flow control. Network addressing concerns the methods whereby a program declares its identity for the purpose of receiving network messages. Flow control involves the various techniques used by DECnet to ensure an orderly transfer of message and control data between nodes in a network.

6.1 Network Addressing

When a program declares its intent to use the network message services, it is influenced by the following DECnet/E features:

- The options available in the Declare Receiver call
- The way in which DECnet/E NSP handles incoming requests for logical link connections
- The multiple copy execution feature, by which more than one copy of a network program can be executed in the RSTS/E timesharing environment
- The automatic job startup feature, by which DECnet/E allows programs to be defined by the system manager for automatic startup on receiving a request for a logical link connection

6.1.1 Declaring Network Names and Object Type Codes

Every DECnet/E program that wishes to receive messages from local senders or establish logical links for data exchange with other network programs must declare its identity in a Declare Receiver call. The program can declare its identity with a nonzero object type and with a null name — abbreviated to DR(n,null) in this chapter — with a zero object type and a nonnull name — DR(0,name) — or with a nonzero object type code and a nonnull name — DR(n,name). Remember from the discussion in Chapter 5 that nonzero object type codes should be used consistently throughout the network.

Declaring identity by object type code alone — DR(n,null) — permits multiple copies of a program to be executed without the necessity of generating a unique identity for each copy. DECnet/E allows more than one program to run and declare identity with the same object type code. Thus, several copies of a program can be executed simultaneously in response to incoming connect requests.

On the other hand, declaring identity by name alone — DR(0,name) — requires uniqueness in the sense that no other program at the local node can declare identity with the same name at the same time. If a program declares its identity with the same name each time it runs, multiple copies cannot execute simultaneously.

It is possible, however, to combine the features of uniqueness and of multiple copy execution. By appending the RSTS/E job number to a common root, a program can create a unique, nonnull name — for example, FAL21, FAL05. (The RSTS/E job number is not directly accessible to a COBOL program but can be determined by executing a MACRO-11 subprogram that uses the UU.SYS directive, described in the *RSTS/E System Directives Manual*.) The program can then declare its identity with both the name created *and* a nonzero object type code. This permits multiple copies of the program to be run as a result of incoming connect requests but permits each copy to create its own unique identity.

6.1.2 Handling of Incoming Connect Requests

When a Connect Initiate Message is received at the local node, NSP searches its list of Receiver ID Blocks for the proper receiver. (All other incoming messages reference an established logical link and will be queued to the correct program without searching the list of Receiver ID Blocks.) The method of addressing specified in the Connect Initiate Message determines how NSP processes the request.

Incoming requests for logical link connections can be addressed to either an object type code alone — CI(n,null) — or to a name alone — CI(0,name). DECnet/E does not allow incoming connect requests to be addressed using both object type code and name. It will reject such requests, sending back a Connect Reject Message with a reason code indicating that the requested program cannot be found.

Incoming Connect Request Addressed to Object Type Code — CI(n,null)

If the incoming Connect Initiate Message is addressed to an object type code — CI(n,null) — NSP checks its Receiver ID Blocks to see if one or more programs are running that have declared identity with this object type code. The Declare Receiver call could have specified either DR(n,null) or DR(n,name).

If one or more such receivers are found that have not reached their declared logical link maximum or message maximum, NSP queues the connect request for the first such receiver in its list (ordered by job number).

NOTE

Remember from Chapter 5 that a program sets its own limits on the number of logical links allowed and the size of the pending message queue in the Declare Receiver call.

If there is no such receiver or if there are one or more such receivers but all have reached their declared logical link maximum or message maximum, NSP attempts to start another copy of the program. This process is described fully in Section 6.1.4 and also in the *DECnet/E System Manager's Guide*. Briefly, NSP checks to see if the program identified in the Connect Initiate Message has been defined by the system manager for automatic startup. If so, NSP creates a job and starts the program. The program must then issue a Declare Receiver call with the appropriate object type code. This creates a Receiver ID Block and a pending message queue to which the waiting Connect Initiate Message can then be transferred. If the program does not issue a Declare Receiver call within two minutes, NSP rejects the link by returning a Connect Reject Message to the remote program.

Incoming Connect Request Addressed to Name — CI(0,name)

When the incoming Connect Initiate Message is addressed to a name — CI(0,name) — NSP checks its list of Receiver ID Blocks to see if a program is running that declared its identity with this name. The Declare Receiver call could have specified either DR(0,name) or DR(n,name).

If such a receiver program is running, and it has not yet reached its declared logical link maximum or message maximum, the Connect Initiate Message is placed in the receiver's queue of pending messages.

If no such program is running, NSP will attempt to start the program. If it is successful in doing so, the program started must declare its identity with the appropriate name. Again, if the program fails to issue a Declare Receiver call within two minutes, NSP rejects the connection.

If such a program is running, but it has reached its declared logical link maximum or message maximum, NSP rejects the connection, returning a Connect Reject Message with an appropriate reason code explaining the rejection. Automatic startup is not attempted in this case, because a program is already running with the name to which the Connect Initiate Message was addressed. A second copy with the same name cannot be generated since receiver names must be unique.

6.1.3 Multiple Copies

In a DECnet/E system, multiple copies of a program can be started by a local terminal user, through the RSTS/E BATCH processor, or through the automatic job startup feature. The manner in which copies of a program can be started has some bearing on the way in which a Declare Receiver call should be issued.

Multiple Copies Initiated by Terminal Users

A program run from a local terminal with the purpose of communicating with a remote program will normally send at least the initial Connect Initiate Message requesting a logical link. (It is unlikely that remote users would wish to have the success of incoming link connections rely upon someone starting the program manually.)

In this case, the program can declare its identity with DR(n,null), or with DR(n,name) or DR(0,name) as long as the program generates a unique name with each copy, as described in Section 6.1.1. The choice depends on whether or not the program receives any connect requests, and if so, if it is necessary for the remote program to connect with a specific copy.

Suppose, for example, that the local program issues the initial Connect Initiate Message for establishing the first logical link. Suppose further that it is then necessary for the remote program to establish a second logical link with the same copy that sent the Connect Initiate Message for the first link. The local program can declare its identity with either the form DR(0,name) or DR(n,name). The unique copy name is passed to the remote program in the Connect Data Block in the Connect Initiate Message. The remote program can then address its connect request in the form CI(0,name) to ensure that the request is delivered to the same copy that issued the first Connect Initiate Message.

Multiple Copies Executed under BATCH

Multiple copies of a program can also be executed through use of the RSTS/E BATCH processor (see the *RSTS/E System User's Guide*). Again, such programs will normally be responsible for initiating at least the first logical link connection. The techniques described previously for multiple copies started by terminal users would also apply here.

Multiple Copies Initiated through Automatic Startup

Multiple copies of a program can also be initiated for programs that the system manager has designated for automatic job startup. Such copies result from incoming connection requests addressed CI(n,null). (See Section 6.1.4 for full details.) The local program can declare its identity with DR(n,null), or with DR(n,name) as long as the name is unique for each copy.

6.1.4 Automatic Job Startup

Using the Network Control Program (NCP), the system manager can install user programs for automatic startup. Such programs reside as files on disk in the normal fashion and can be started by NSP to fulfill incoming requests for logical link connections (see Section 6.1.2). This installation method is described fully in the *DECnet/E System Manager's Guide*. Briefly, a DEFINE OBJECT or SET OBJECT command in NCP relates the file name to, again, an object type code, a name, or both. These commands also allow two parameters to be associated with a program. The second parameter is made available to the program when it is started automatically. (The first parameter is used as a starting line number for BASIC-PLUS or BASIC-PLUS-2 programs. It is ignored for COBOL programs.)

NSP maintains a list of these definitions in its parameter file. The definitions are ordered by object type code, and thus, there can be only one program defined for automatic startup for each object type code. In particular, there can be only one program defined for automatic startup with an object type code of zero — (0,name).

When a program must be started automatically, NSP checks the object type code or name given in the incoming connect request against its list of defined programs. If a match is found, NSP executes a monitor request to run the program as a detached job under RSTS/E.

Once started, the program must declare its identity with the object type code or name that was specified in the Connect Initiate Message that caused the program to be started. If it does not, the connect request will never be delivered to the proper pending message queue.

NSP passes this information to the started program in core common. Any optional parameter established by the DEFINE OBJECT or SET OBJECT command is also passed in the string. A COBOL program cannot access core common directly, but a MACRO subprogram can be used to retrieve the data. (See the *RSTS/E System Directives Manual*.) The format of the core common string is as follows:

| Byte(s) | Contents |
|----------|--|
| CORCMN+0 | Length of data in core common |
| +1 | Object type code (0-255) |
| +2-7 | Name (ASCII) |
| +8-9 | Second DEFINE parameter (-32768 through 32767) |
| +10 | Reserved |

The mnemonic CORCMN is assigned by the file COMMON.MAC to octal location 460 (see the *RSTS/E System Directives Manual*).

If the object type code in byte 1 of core common is nonzero, the program must declare its identity with that object type code. It can also use a nonnull unique name, if desired. If the object type code is zero, the program must declare its identity with the name given in bytes 2-7 of core common. It can also declare its identity with a nonzero object type, if desired.

When the Declare Receiver call is issued, NSP transfers the waiting connect request to the program's queue of pending messages as a Connect Initiate Message. NSP waits two minutes from the time it starts the program for the connect request to be transferred.

If the connect request is not transferred to a queue within this time period, NSP rejects the connection, sending a Connect Reject Message to the remote program. Since the remote program could also abort the link within this period, the first action of an automatically started program (after issuing the Declare Receiver call) should be to issue a Receive call to check for pending messages. If there are none, the program should issue a Remove Receiver call, and terminate execution with the Kill Job system call. (This call is not directly accessible to a COBOL program but can be executed through a MACRO subprogram using the UU.CHU system call. See the *RSTS/E Sys-*

tem Directives Manual). Once the connect request has been received, the program should respond to the request immediately. While the local NSP's timer is stopped when the message is transferred to the receiver's pending message queue, the remote node could be executing a timeout on the Connect Initiate Message.

A program that has been started in response to an incoming Connect Initiate Message runs as a privileged RSTS/E user, and it is the responsibility of that program to protect the rest of the system from unauthorized outside access. Thus, after receiving the Connect Initiate Message, the program should check the accounting information passed in the Connect Data Block to determine the user account (and, therefore, the privilege level) in which it should run. If no such information is specified in the Connect Data Block, the program should issue a Get Local Node Parameters call to obtain the local node's default project-programmer number, if any. If a default account has been specified, the program should then look up the password for that account using the Read Accounting Data system call and log into the default account. (These system functions are not directly accessible to a COBOL program but can be executed through MACRO subprograms using the UU.RAD and UU.LIN system calls. See the *RSTS/E System Directives Manual*). If the local node has no default account, the connect request should be rejected. (Programs that always log into a predetermined account need not go through this procedure.)

6.1.5 Examples of Network Addressing

This section presents three examples of network addressing using the features just discussed.

The first example (Figure 6-1) shows the network addressing used by two DECnet/E utility routines — the Network File Transfer utility (NFT) and the File Access Listener (FAL). NFT is run by local users to connect with a remote FAL to access files at remote nodes. FAL is run at the local node as a result of incoming connect requests from a remote NFT. Both programs are designed to execute multiple copies as demand requires.

Since NFT is started only by local terminal users and never by remote request, an object type code is not necessary. NFT is loaded from the system disk and declares its identity with an object type code of zero and a unique name each time it is run. NFT issues the Connect Initiate Message to establish a logical link for the transfer of data to or from remote nodes as specified by the user.

Copies of FAL are started by incoming Connect Initiate Messages addressed to object type code 17. FAL declares its identity with object type code 17 and a unique name. It also declares that the maximum number of logical links allowed is one so each incoming connect request starts a new copy.

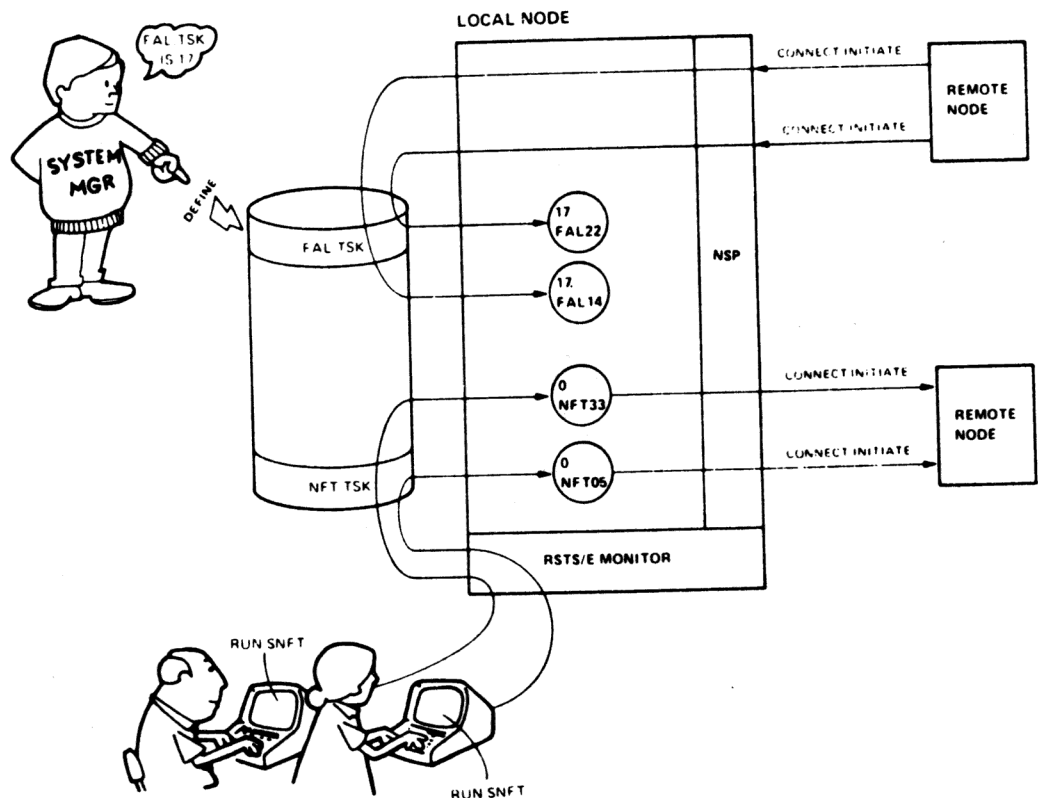


Figure 6-1: Network Addressing Used by NFT and FAL

The second example (Figure 6-2) shows NETCPY, a device copying program that can be started either by local terminal users or by incoming connect requests. The program copies an entire device between nodes. Multiple copies are executed according to demand. Here, the Declare Receiver call is different depending on how the program is started. NETCPY is a BASIC-PLUS program and takes advantage of the starting line number feature. Using the Network Command Program (NCP), the system manager issues a command to SET or DEFINE NETCPY.BAC as object 20, with a starting line number of 29000. If NETCPY is started by a Connect Initiate Message from a remote

program addressed to object type code 20, the program declares its identity with object type code 20, a unique name, and a link maximum of one. If it is started by a local terminal user, it declares its identity with an object type code of 0, a unique name, and a link maximum of zero. Declaring a zero link maximum ensures that a program started by a local user will never have a Connect Initiate Message queued from a remote program.

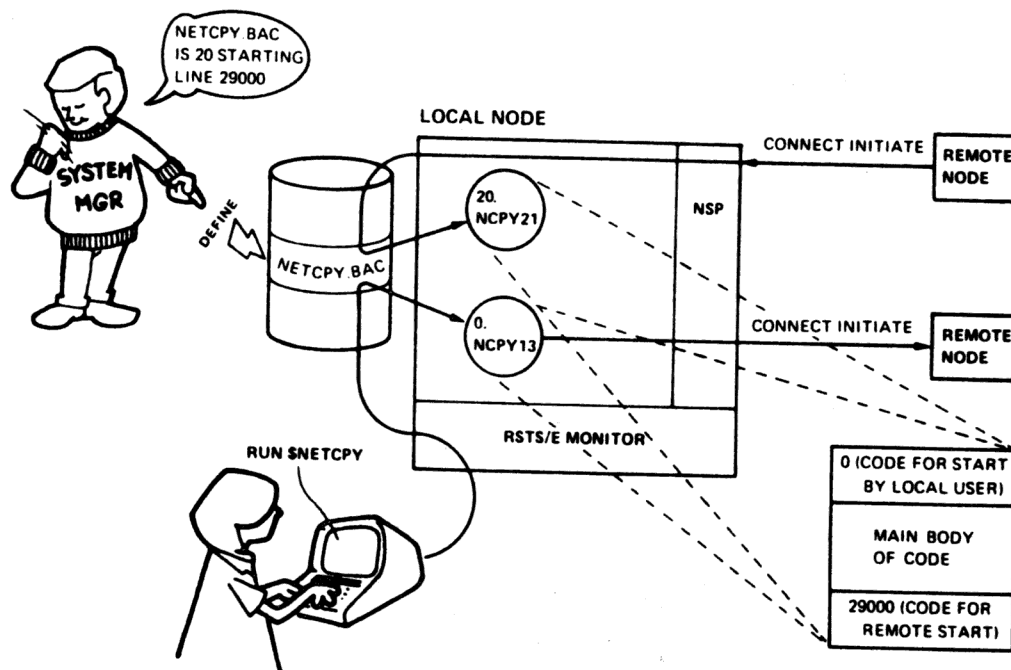


Figure 6-2: Network Addressing Example — Different Declare Receiver for Local and Remote Start

The third example (Figure 6-3) shows a program that can be run either by a local terminal user or as a result of an incoming connect request. The local terminal user is presented with a data display. Parameters are selected by the user and passed to a remote program monitoring a data gathering device. At the end of the day, the remote program passes the data back to the local program for analysis and to the local terminal user, who repeats the process.

Since only one copy of the program is ever needed at any given time, it declares its identity with a zero object type code and the name STAT. It is stored as file RSTAT.SAV on disk, where it can be run by the local user. Using NCP, the system manager has issued a command to SET or DEFINE RSTAT.SAV as STAT.

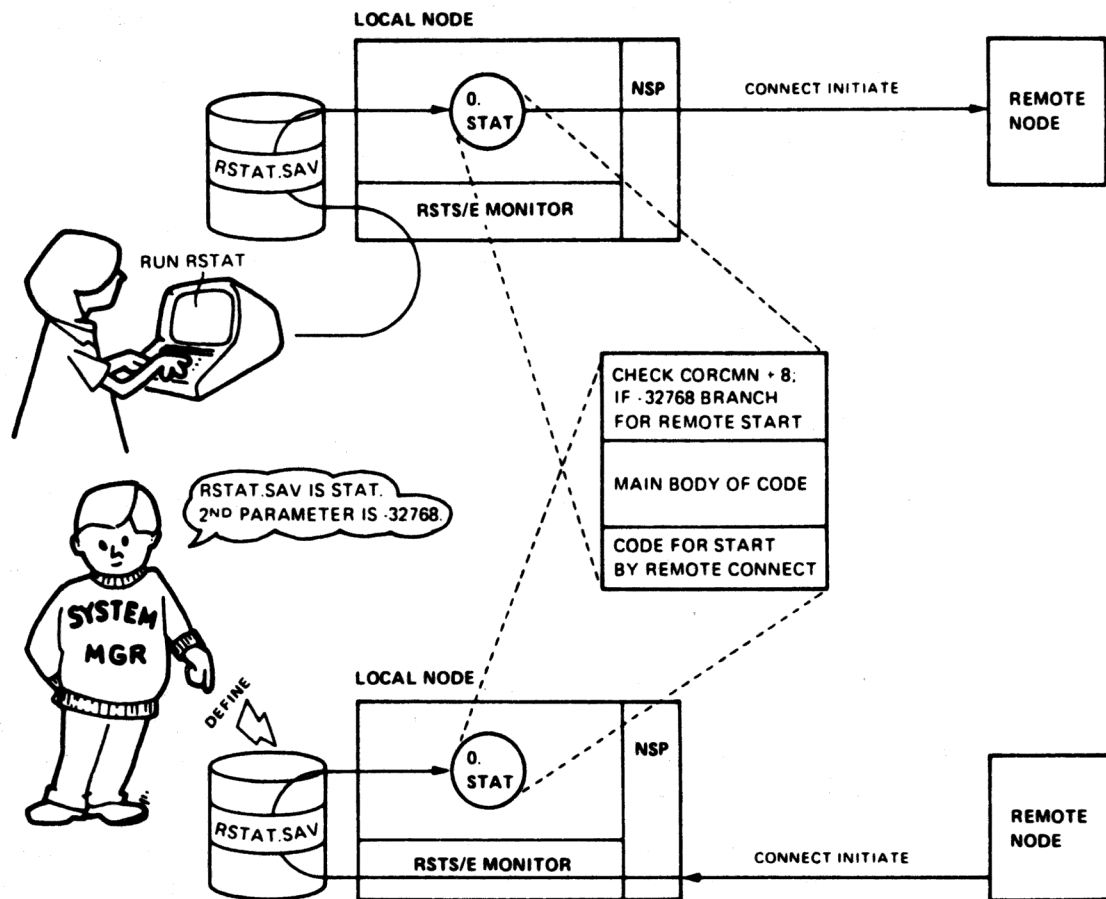


Figure 6-3: Network Addressing Example — Declaring Identity by Name Alone

6.2 Flow Control

DECnet flow control protects user programs as well as entire systems from being flooded with more messages than can be processed at any given time.

A DECnet/E sender is subject to flow control regulations imposed by the local NSP, the remote NSP, and the remote program (if the remote program requested flow control when the link was established). As a receiver, a DECnet/E program can select from three flow control options for a logical link to regulate the amount of data sent to it by the remote program.

As a sender and/or receiver, a program must deal with the following four aspects of DECnet/E operation:

- **Transmit queue management**, imposed by the local (sending) NSP, ensures that a local program does not use too much system buffer space for transmission.
- **Backpressure flow control**, imposed by the remote NSP, is used as a safeguard against excessive incoming data messages.
- **Data Message flow control**, imposed by receiving user programs, enables a program to control the flow of data to it. Both programs using a logical link can (independently) choose to receive network messages only when they specifically request data from the other program; or each can choose no flow control, placing no restrictions on the amount of data that they will receive.
- **Interrupt Message flow control**, imposed by either the remote NSP or the remote program (depending on the DECnet implementation), protects against Interrupt Message congestion. A DECnet/E program can receive only one unsolicited Interrupt Message over a link. After that, the program must request an Interrupt Message before the remote program can send one.

6.2.1 Transmit Queue Management

Once a logical link is established, DECnet/E maintains two transmit queues for the link: a transmit queue for Network Data Messages and a transmit queue for Interrupt and Link Service Messages. Both queues hold messages waiting for acknowledgment from the remote NSP.

For example, when a DECnet/E program sends a Network Data Message to a remote program over a logical link, the message is queued to the Transport software for transmission and is also placed in the data transmit queue. The Network Data Message is held there until the local NSP receives an acknowledgment from the remote NSP indicating that it has received the message.

If NSP does not receive an acknowledgment within a reasonable period of time, it assumes that the message has been lost and requeues it to Transport for retransmission. If the message has not been successfully transmitted after a specific number of retries, NSP aborts the link with a reason code indicating that the remote system is not responding.

The timeout period used depends on the message type. For outgoing Connect Initiate Messages, a predetermined value is used. Once a logical link is established, however, network management dynamically computes the timeout period using various network parameters defined by the system manager using NCP and the SET EXECUTOR command. (See the *DECnet/E System Manager's Guide* for details.)

This acknowledgment procedure between the two NSPs ensures that messages sent over the logical link have been received, and that they were received in the same order they were sent.

Each of the transmit queues for a logical link can hold up to eight messages. (This limit can be changed by the system manager to any number from one to eight to further protect system buffer space.) When one of the transmit queues for a logical link is filled with messages waiting for acknowledgment, further transmission of that type of message over the logical link is inhibited by the local NSP until one or more of the queued messages have been acknowledged by the remote NSP.

An error is returned to the program if it issues a send request while the transmit queue for that message type is full. This condition is temporary. When the condition clears, send requests are again allowed for that message type. A program getting such an error on a send request should simply wait a short time (1 to 5 seconds) and retry the send request.

The state of the transmit queues also affects the sending of Disconnect Messages. If any messages are waiting for acknowledgment in the Data or Interrupt/Link Service transmit queue, a Send Disconnect Message call for that logical link will not be allowed and will result in an error. Thus, a program sending a successful Disconnect Message can assume that all previously sent messages have been received and acknowledged by the remote NSP. A Link Abort Message for that logical link will be allowed, however, and any messages waiting in the transmit queues will be discarded.

Figure 6-4 illustrates transmit queue flow control. PROGA has established two logical links, and PROGB and PROGC have each established one logical link. Each logical link has its own Data and Interrupt/Link Service Message transmit queues. PROGA issues a Send Network Data Message system call. The message goes to the Data transmit queue for that logical link and is also queued for transmission by the communication hardware. At the device driver level, all network messages (both data and control messages) await transmission over the physical line.

PROGB is shown trying to send a Network Interrupt Message. The Interrupt/Link Service transmit queue for that link is full, however, so NSP returns an error on the send. PROGC issues a Disconnect Message for its logical link. Since messages are still waiting for acknowledgment in the transmit queues, the Disconnect Message is not allowed and an error is returned. (Sending a Link Abort Message would have broken the link but the messages would have been discarded.)

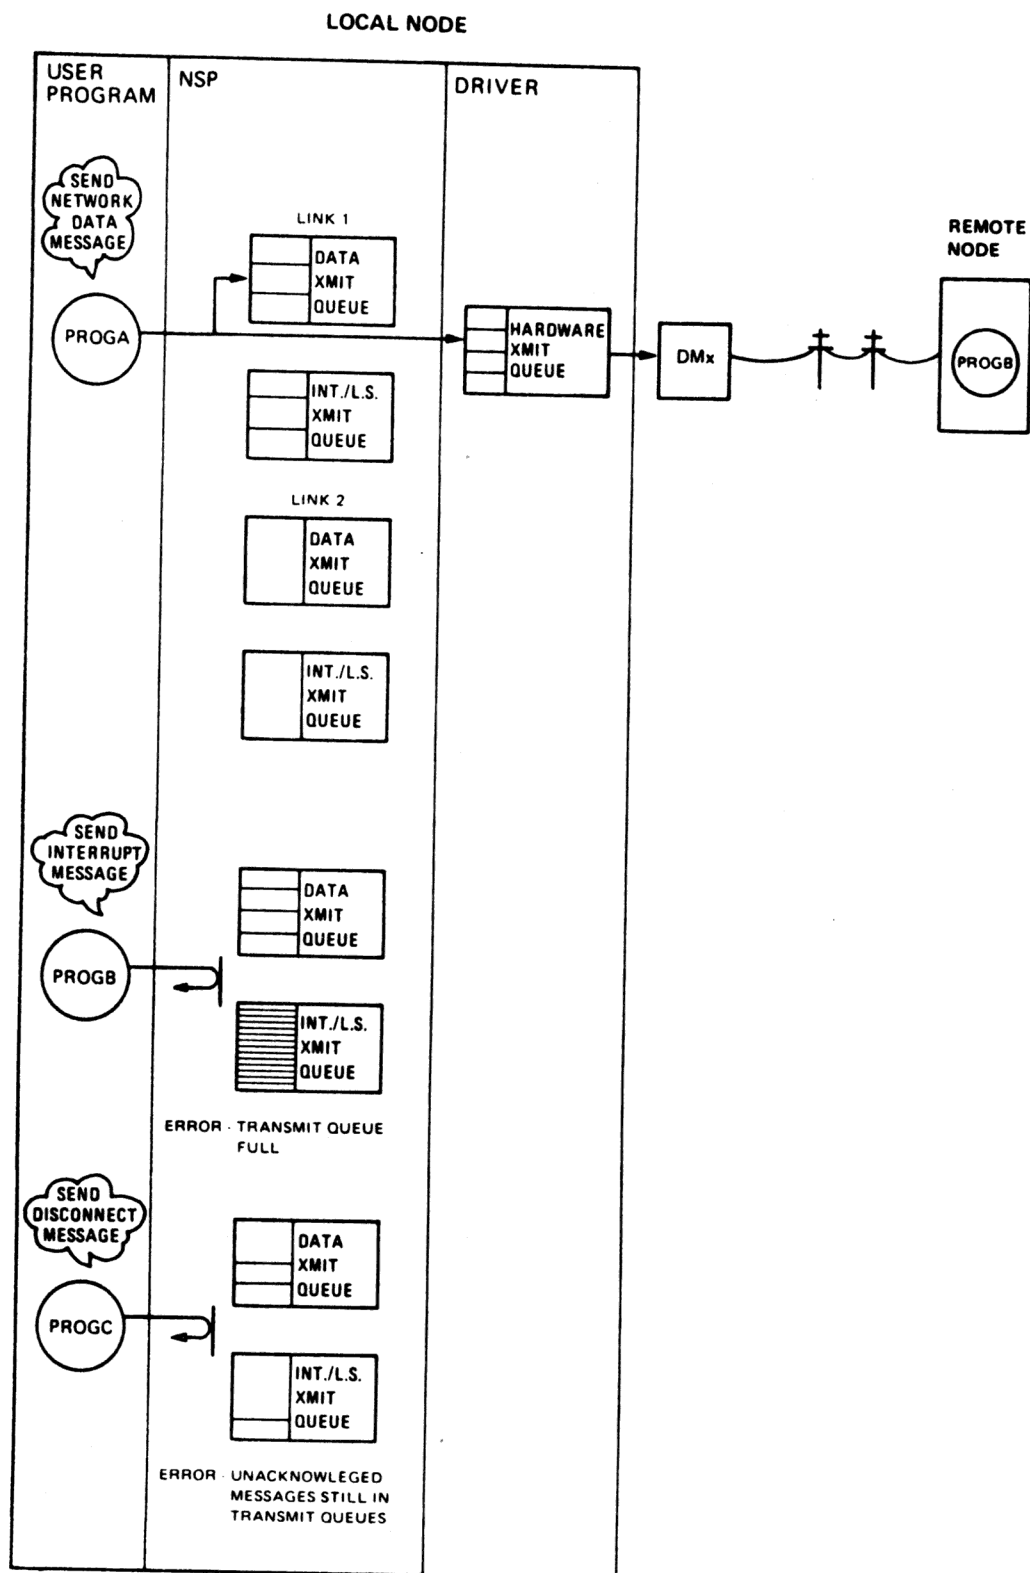


Figure 6-4: How Transmit Queue Flow Control Affects a DECnet/E Program

6.2.2 Backpressure Flow Control

NSP at the remote node can inhibit the local program from sending Network Data Messages over a logical link. The reasons for this blocking can vary from system to system: lack of system buffer space, system load, or some similar condition. This type of flow control, whereby the *remote* NSP can inhibit message flow from the local program, is known as backpressure.

When the remote NSP receives a Network Data Message that it cannot process, it returns a negative acknowledgment (NAK) for the message, along with an error code informing the local NSP that data message flow to the remote program is being inhibited on the link over which the message was sent. The local NSP then inhibits transmission of Network Data Messages over the indicated logical link until further notice is received from the remote NSP. (Note that the local program can still receive messages over the logical link.) When the remote NSP informs the local NSP that the backpressure condition has cleared, the local NSP retransmits the message that was previously NAKed and reenables transmission of Network Data Messages over the logical link.

If the backpressure condition clears before the local program attempts another send over the link, the program is not informed that the condition occurred. However, if the program attempts to send another Network Data Message while transmission on the link is still inhibited, the local NSP returns an error to the program (see A, in Figure 6-5). When the condition clears, the local NSP delivers a Link Service Message to the program, informing it that transmission of Network Data Messages is again allowed for the indicated logical link (see B, in Figure 6-5). This Link Service Message is delivered to the local program on a Receive call when the queue of pending messages is empty.

Note that the remote NSP enforces flow control for specific logical links, not for the entire node.

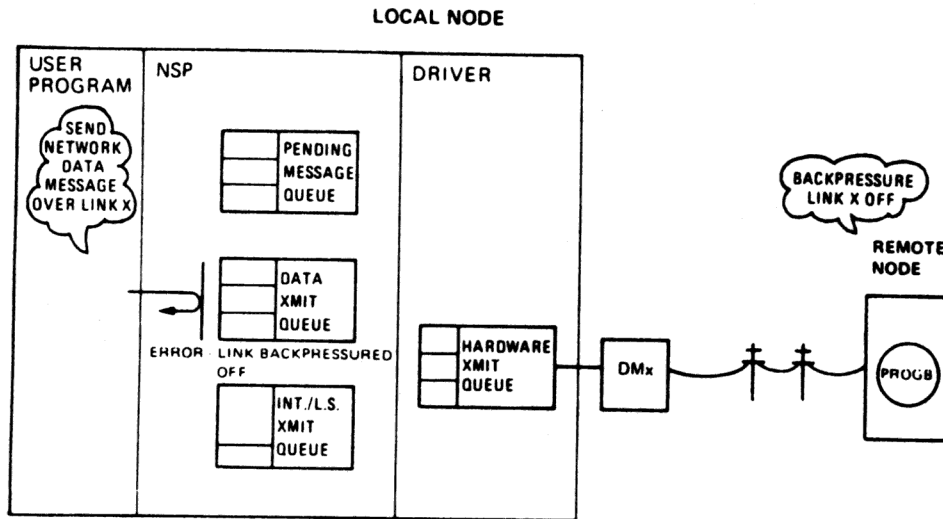
From the reverse viewpoint, the local NSP will inhibit incoming data messages on a logical link when a Network Data Message is received from a remote program over the link and the pending message queue of the intended receiver is full. (Remember that the local program itself limits the size of its pending message queue in its Declare Receiver call.) The local NSP informs the remote NSP that incoming data message flow has been inhibited by sending a message to the remote system.

Under these conditions, the Network Data Message is not queued for the local program, but is NAKed by the local NSP. This forces the remote NSP to retransmit the message after flow is reenabled. Any data messages already transmitted by the remote node will also be NAKed by the local NSP.

Furthermore, the local NSP marks all other logical links on the Receiver ID Block to be turned off if any Network Data Messages are received on those links. However, no message is sent to the remote system unless a Network Data Message is received on one of these links before the restriction is lifted. (Note that while the local NSP will no longer accept incoming messages, the local program can still transmit.)

When the local program empties its pending message queue, any links that were actually turned off are turned back on and a message is sent to the remote NSP to reenale transmission. Those links that were merely flagged to be turned off are restored to normal operation, but no message is sent to the remote system.

A. Remote node backpressures a link off



B. Remote node clears backpressured link, local program notified on receive

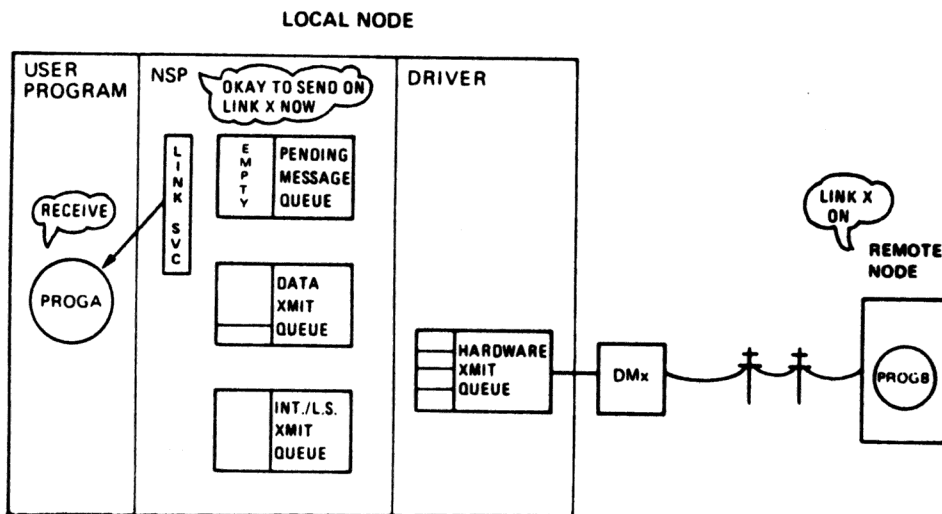


Figure 6-5: How Backpressure Flow Control Affects a DECnet/E Program as a Transmitter

6.2.3 Program-Regulated Flow Control

During the connection sequence that establishes a logical link, each program must select one of the following three flow control options to regulate incoming data over the link:

- **Segment flow control.** The program requests data from the remote program by indicating how many segments it will accept at any given time. A segment is defined as the amount of user data transferred in one Network Data Message. The maximum size of a segment is determined by the receive maximum established in the Connect Initiate or Connect Confirm Message. The sending program is not required to use this maximum, however, and can, in fact, send smaller segments. Note also that segments transmitted and received by the same program are not necessarily equal in size.
- **Logical message flow control.** The program requests data from the remote program by indicating how many logical messages it will accept at any given time. A logical message can consist of one or more data segments. Control flags, included as part of a Network Data Message, indicate whether the data segment is the beginning, middle, end, or sole segment of a logical message.
- **No active flow control.** The program sets no limits on the rate at which messages can be sent to it over a logical link.

If the program selects segment flow control or logical message flow control, the remote NSP will not send Network Data Messages over the link until the local program requests such transmission. Requests are made by sending a Link Service Message requesting a specific number of segments or logical messages.

If the program selects the third option, no active flow control, no Link Service Messages can be sent except those requesting Interrupt Messages. Network Data Messages will be transmitted by the remote NSP as the remote program sends them, subject to backpressure.

Flow control selection is independent at both ends of the logical link. That is, the remote and local programs need not select the same type of flow control. The option selected by the remote program, however, affects how the local program can send Network Data Messages. Therefore, the following discussion of flow control is presented from two viewpoints:

1. How the option selected by the local program affects incoming Network Data Messages
2. How the option selected by the remote program affects the way the local program sends Network Data Messages

Flow Control on Incoming Data

If the local program selects logical message flow control in the Connect Initiate or Connect Confirm Message establishing a logical link, the local program must request data from the remote program by sending Link Service Messages specifying the number of logical messages to be transmitted.

When the Send Link Service Message call is issued, the local NSP adds the value specified to a message counter maintained for the link. The local NSP also sends the information on to the remote NSP. The remote NSP likewise adds the value specified to a message counter it maintains for the link. Thus both NSPs count the outstanding message requests. Each time the remote NSP sends the end segment of a logical message, it decrements the message counter. When the counter reaches zero, the remote NSP will inhibit transmission until the local program sends another Link Service Message specifying another value to add to the message counter.

The local NSP uses its message counter as a protective measure. Each time it receives a data segment with the end-of-logical-message flag set, it decrements the message counter. When the counter reaches zero, the local NSP no longer accepts data segments over the link. Should the remote NSP send a segment while the local NSP's counter is zero, the message is discarded and the link is aborted. The local NSP places a Link Abort Message in the local program's queue of pending messages and sends a Link Abort Message to the remote NSP. (This situation is a protocol violation and should not occur with correct NSP operation.)

When the local program selects segment flow control in the Connect Initiate or Connect Confirm Message, the local and remote NSP maintain segment counters for the logical link. The remote program can still send segments with the logical message control flags set for the local program to process internally, but the counters maintained by NSP are kept for data segments sent, not for logical messages. Data is requested by sending a Link Service Message indicating the maximum number of segments to be sent. The local and remote NSPs add this value to their data segment counters and then decrement the counters as the remote program sends segments. When the data segment counter reaches zero, transmission over the logical link is suspended until another Link Service Message increments the count.

When no active flow control is requested in the Connect Initiate or Connect Confirm Message, segments from the remote program can be transmitted over the logical link without being requested by the local program. The local and

remote NSPs do not keep counters to regulate data sent and Link Service calls are not used to request transmission of data. Data Message flow can still be restricted by backpressure, however.

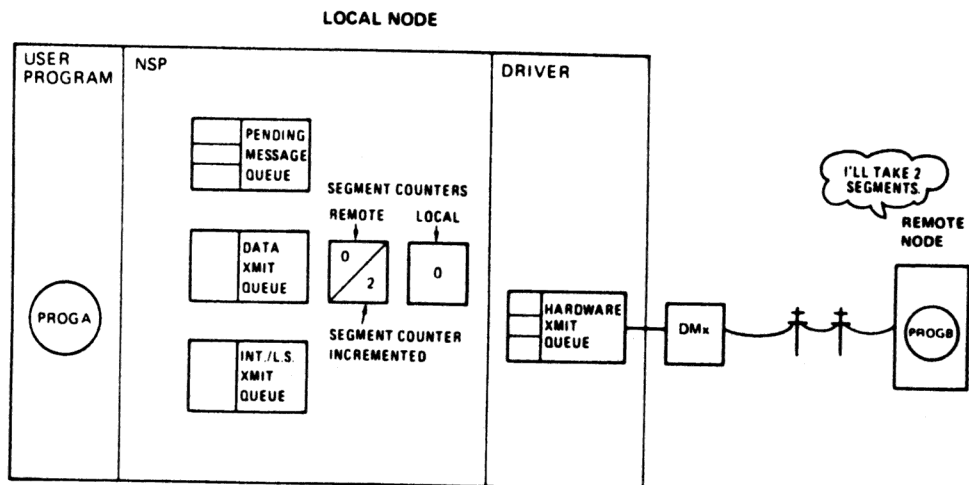
Flow Control on Outgoing Data

When the remote program specifies a flow control option in its Connect Initiate or Connect Confirm Message establishing the logical link, the local NSP notes the option and establishes the appropriate counter. The counter is increased when a request for data is received from the remote system or program. If the remote program requested segment flow control, the counter is decremented by one each time the local program issues a Send Network Data Message call. If the remote program requested logical message flow control, the message counter is decremented by one each time a Network Data Message is sent that is flagged as the end data segment of a logical message.

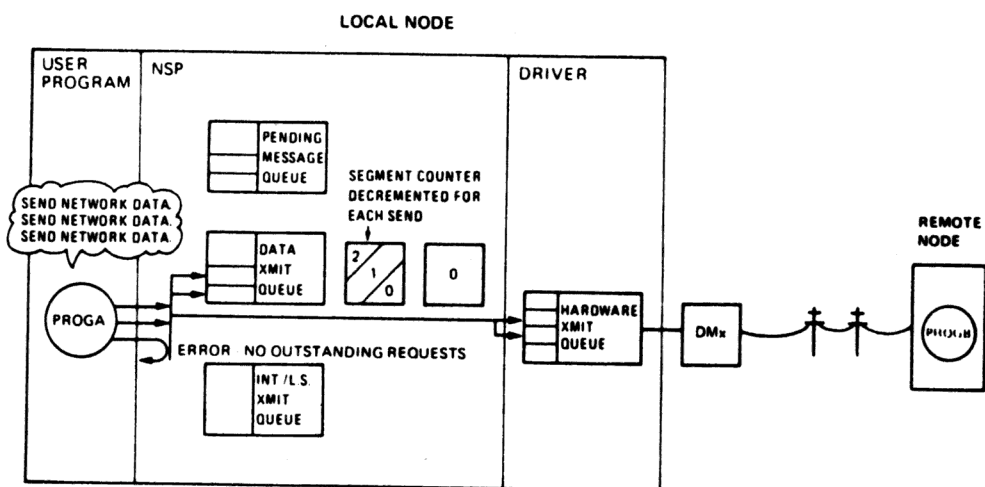
When the counter equals zero, the local NSP inhibits further transmission of Network Data Messages over the logical link, returning an error if the local program tries to do so. In this circumstance, the local NSP delivers a Link Service Message to the local program when the condition clears and the pending message queue is empty. A Link Service Message is not delivered, however, if the local program did not attempt to send a Network Data Message while the request count was zero.

This process is illustrated in Figure 6-6. The remote program has requested segment flow control. First, the remote program requests two segments (see A, in Figure 6-6). The segment counter for the remote program is incremented, but since PROGA has not had transmission inhibited, no Link Service Message is queued. The local program issues three Send Network Data Message calls (see B, in Figure 6-6). The segment counter is decremented to one with the first send, and to zero with the second send. The third send, while the counter is zero, results in an error. When the remote program receives and processes its two segments, it asks for two more segments (see C, in Figure 6-6). The segment counter is incremented from 0 to 2, reenabling data sends. Since PROGA attempted a send that failed, a Link Service Message is delivered to PROGA to indicate that data flow over the link has been reenabled. The Link Service Message is delivered to PROGA when it executes a Receive call when the pending message queue is empty.

A. Remote program requests two segments



B. Local program sends two segments, third send returns an error



C. Remote program requests two segments, local program notified on receive

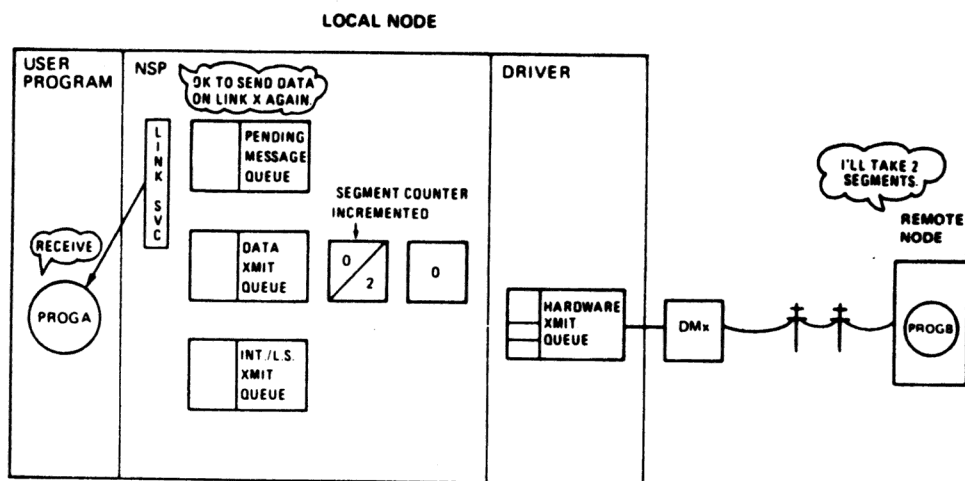


Figure 6-6: How the Flow Control Option Selected by the Remote Program Affects a DECnet/E Program as a Transmitter

6.2.4 Interrupt Message Flow Control

The transmission of Interrupt Messages over a logical link is subject to flow control by a request count mechanism. This mechanism is controlled by either the remote NSP or the remote program, depending on the particular DECnet implementation.

All logical links are initiated with an outstanding Interrupt Message request count of one from the remote system. If the local program sends an Interrupt Message over the logical link, this count is decremented to zero. The local NSP does not allow another Interrupt Message to be sent over this logical link until the remote NSP or remote program sends a Link Service Message, incrementing the request count. An attempt by the local program to send an Interrupt Message while the counter is zero will result in an error.

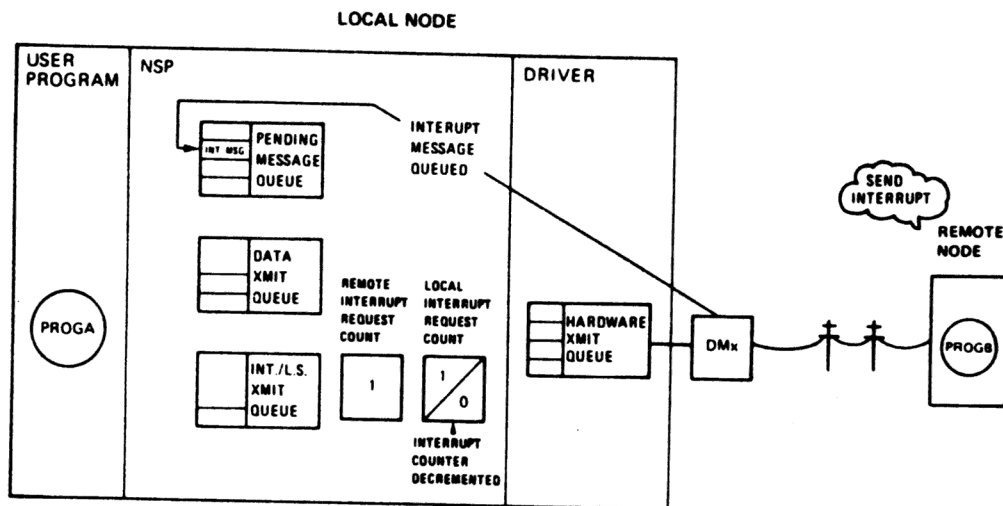
Once the request count is incremented by the remote system or program, the local NSP reenables Interrupt Message flow. If the local program has tried to send an Interrupt Message while the count was zero, NSP notifies the local program with a Link Service Message that Interrupt Messages can again be sent over this logical link. This Link Service Message is delivered to the local program on a Receive call when the pending message queue is empty. If the local program has not tried to send an Interrupt Message over the link while flow was disabled, no Link Service Message is delivered.

From the reverse viewpoint, all DECnet/E logical links are initiated with an outstanding request for one Interrupt Message. If the remote program sends an Interrupt Message over the link, the local NSP decrements this count to zero. No further Interrupt Messages are accepted until the local program requests one by sending a Link Service Message to increment the count. This also increments the request counter at the remote end and thus allows the remote program to send an Interrupt Message if it so desires. DECnet/E NSP does not allow the local program to request an Interrupt Message for a logical link unless a previous Interrupt Message on the link has been received.

In DECnet/E, only one Interrupt Message can be requested on a logical link at any one time. That is, the request count is either zero or one. If a program tries to issue a Link Service Message requesting an Interrupt Message on a link when the request count is already one, the Link Service Message will result in an error.

Figure 6-7 illustrates Interrupt Message flow control for a DECnet/E receiver. The remote program sends an Interrupt Message over a logical link (see A, in Figure 6-7). The local NSP decrements the Interrupt Message counter to zero and does not accept further Interrupt Messages over this link. The remote NSP also decrements its own counter and inhibits further Interrupt Messages from being sent. The local NSP maintains its counter as a safeguard measure. The local program sends a Link Service Message requesting one Interrupt Message over this logical link (see B, in Figure 6-7). The local NSP increments its counter and tells the remote NSP to reenables Interrupt Messages over the link. The local program must receive the first Interrupt Message before the Link Service Message can be sent.

A. Remote program sends an Interrupt Message over a logical link



B. Local program reenables interrupts for the link (allowed only when previous Interrupt Message for this logical link has been received)

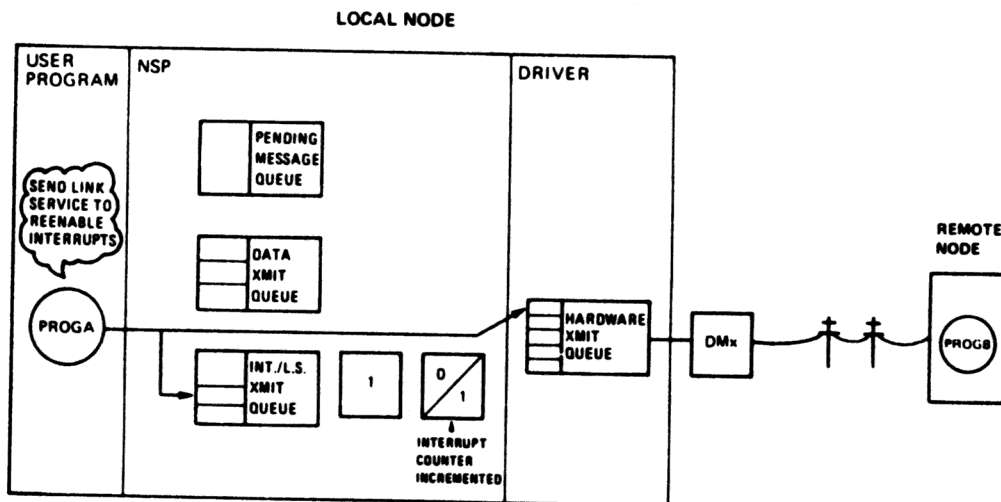


Figure 6-7: How Interrupt Message Flow Control Affects a DECnet/E Program as a Receiver

6.2.5 Programming Hints for Flow Control

A DECnet/E program can be prevented from transmitting for a number of reasons:

1. No buffer space is available at the local node. (Prohibits sending of all message types.)
2. The Data transmit queue is full. (Prohibits sending of Data Messages.)
3. The Interrupt/Link Service transmit queue is full. (Prohibits sending of Interrupt and Link Service Messages.)

4. Message flow over the link has been inhibited by the remote system because of a backpressure condition. (Prohibits sending of Data Messages.)
5. No requests are outstanding for segments or logical messages from the remote program. (Prohibits sending of Data Messages.)
6. No requests are outstanding for Interrupt Messages. (Prohibits sending of Interrupt Messages.)

The first three conditions indicate congestion at the local node. They are temporary and no notification is given to the program when the condition clears. When one of these errors occurs on a send call, the program should simply reissue the call after a short delay of one to five seconds.

The last three conditions, however, indicate congestion at the remote node. Since there is no way to predict how long it can take for one of these conditions to clear, the local NSP does not allow further send requests of Data or Interrupt Messages until the remote system indicates that the condition has cleared.

When sending is blocked for one of these last three conditions, any attempt to send a message by the local program results in an error. The program can still receive messages, however. If an attempt to send a message resulted in an error in this fashion, NSP notifies the program when the condition clears. It does this by delivering a Link Service Message to the program when the program issues a Receive call and the pending message queue is empty. No notification is given if the local program did not try to send a message.

Thus a program whose main function is to transmit data can be designed to send data over a logical link until transmission is blocked by one of the last three conditions. When sending is inhibited, the program can issue Receive calls to process all pending messages. When the pending message queue is empty and all conditions that block flow have cleared, a Link Service Message is delivered on a Receive call, giving current status information on flow control for the logical link. (If more than one logical link has been blocked, Receive calls should be issued for all links until no messages are returned, since link service information can also be available on the other links.)

A program whose main function is to receive data is concerned with keeping the pending message queue emptied so that the queue does not fill and cause the local NSP to inhibit incoming message flow over the program's logical links. The Receive call with sleep can be used effectively by such a program so that it will be notified when a message is delivered to the queue. In this case, a second Receive call must be issued to retrieve the message.

An interactive program can choose between these mostly-send and mostly-receive approaches, according to the requirements of the application. A program can also take advantage of the link status information that can be returned after sending Network Data, Interrupt, and Link Service Messages to regulate send and receive requests. The link status information returned is in the same format as a received Link Service Message.



Chapter 7

Extended COBOL Interface Call Formats

This chapter presents the specific calls, with detailed formats, used to access the DECnet/E message services with the extended COBOL interface.

7.1 Programming Background

To use the RSTS/E message services in the extended COBOL interface, you code subprogram calls with the CALL statement, including the USING clause to pass operands to the subprograms. These subprograms reference MACRO-11 subprograms stored in the system library that must be linked with your program. Thus, the process involves coding, compiling, and linking. These aspects of using the extended DECnet/E interface in COBOL are described briefly in the following subsections. Specific call formats are given in Sections 7.2 - 7.7.

7.1.1 Coding

The general format for the extended COBOL message calls is:

CALL "name" USING *arg1*, *arg2*,...

where

name is the name of one of the twelve subprograms available for message services. Each subprogram processes a call at execution time by executing the following steps:

1. Checking for the proper number of arguments in the call.
2. Translating the call and arguments into the form recognizable to the RSTS/E monitor.
3. Passing control to the RSTS/E monitor and, therefore, to NSP. (In DECnet/E systems, NSP is part of the RSTS/E monitor.)

NSP processes the request and returns control to the user program. If an error occurs, NSP returns an error value to an integer variable defined as an argument in each call.

arg... is the argument list for the particular call. Each argument must have been defined as a data item in the Data Division of the program.

Note that the message subprograms are MACRO (assembly language) object code subprograms and the numeric arguments passed and returned require definition in the Data Division of your COBOL program so that they fit in one computer word. Such a one-word COMP data item can be defined with PICTURE statements ranging from

PIC 9(1) COMP through PIC 9(4) COMP (for unsigned values)

to

PIC S9(1) COMP through PIC S9(4) COMP (for signed values)

Once a COMP data item is defined as PIC S9(4), for example, values ranging from -32768 through 32767 can be returned to the data item (even though five character positions are used). To pass a value of more than 4 characters, you MOVE (in the Procedure Division) the value to the data item, rather than use the VALUE IS clause in the Data Division.

Also, with the extended COBOL interface, you must be able to access one-byte binary values. Some general tips on how to do this are given below. Other examples are given in the call format discussions in Sections 7.2 - 7.7.

Tips for Returned Status Information in Byte Form

The Receive call and three send calls (Send Network Data, Interrupt, and Link Service Messages) return status and control information to a 40-byte user buffer. Some of this information consists of signed and unsigned numeric byte values. PDP-11 COBOL recognizes no such data category. Thus, to use this information you must convert them to one-word integer values. The technique is different for signed and unsigned byte values. Two examples follow. The first converts signed one-byte binary items to signed one-word COMP items. If the value expected is within the range -128 through 127, you can use the technique shown in the first example. The second example converts unsigned one-byte binary items (ranging from 0 to 255) to unsigned one-word COMP items. If the value expected is within the range 0 to 255, use this or a similar technique.

Example 1 — Signed

```
IDENTIFICATION DIVISION.  
PROGRAM ID. BYTSGN.  
AUTHOR.  
      R-FRIED.  
DATE-WRITTEN. 18-JUN-79.  
DATE-COMPILED.  
      18-JUN-79.  
  
ENVIRONMENT DIVISION.  
CONFIGURATION SECTION.  
SOURCE-COMPUTER.      PDP-11.  
OBJECT-COMPUTER.      PDP-11.
```

(continued on next page)

```

DATA DIVISION.
WORKING-STORAGE SECTION.
01    WORK-WORD          PIC S9(4) COMP.
01    WW REDEFINES WORK-WORD.
      03 LOW-BYTE PIC X.
      03 HIGH-BYTE PIC X.
01    SEVEN-BITS-G PIC 9(4) COMP VALUE 127.
01    SEVEN-BITS REDEFINES SEVEN-BITS-G PIC X.
01    ALL-BITS-G PIC 9(4) COMP VALUE 255.
01    ALL-BITS REDEFINES ALL-BITS-G PIC X.
LINKAGE SECTION.
01    ONE-BYTE PIC X.
01    ONE-WORD PIC S9(4) COMP.

PROCEDURE DIVISION USING ONE-BYTE ONE-WORD.
SBEGIN.
      IF ONE-BYTE > SEVEN-BITS
        MOVE ALL-BITS TO HIGH-BYTE
      ELSE
        MOVE 0 TO WORK-WORD.
      MOVE ONE-BYTE TO LOW-BYTE.
      MOVE WORK-WORD TO ONE-WORD.
      EXIT PROGRAM.

```

Example 2 — Unsigned

```

IDENTIFICATION DIVISION
PROGRAM-ID. BYTUNS.
AUTHOR.
      R-FRIED.
DATE-WRITTEN. 18-JUN-79.
DATE-COMPILED.
      18-JUN-1979.

ENVIRONMENT DIVISION.
CONFIGURATION SECTION.
SOURCE-COMPUTER.      PDP-11.
OBJECT-COMPUTER.      PDP-11.

DATA DIVISION.
WORKING-STORAGE SECTION.
01    WORK-WORD          PIC 9(4) COMP.
01    WW REDEFINES WORK-WORD.
      03 LOW-BYTE PIC X.
      03 HIGH-BYTE PIC X.
LINKAGE SECTION.
01    ONE-BYTE PIC X.
01    ONE-WORD PIC S9(4) COMP.

PROCEDURE DIVISION USING ONE-BYTE ONE-WORD.
SBEGIN.
      MOVE 0 TO WORK-WORD.
      MOVE ONE-BYTE TO LOW-BYTE.
      MOVE WORK-WORD TO ONE-WORD.
      EXIT PROGRAM.

```

Note that not all the status information returned needs to be used in this way. Some values are returned as signed or unsigned oneword COMP data items that you can reference directly with data-names defined as PIC S9(4) COMP or PIC 9(4) COMP.

Tips for Setting Bits

Some of the argument descriptions in the extended COBOL interface calls require that you set bits in one-word COMP data items. Such arguments are usually described as the sum of n bit values $x1 + x2 + \dots + xn$.

The x_i components are then described in terms of what happens when x_i equals zero or some power of two: 2, 4, 8, ..., 128. Setting a component to the nonzero value sets the bit that has that value in the computer word. Setting a zero value clears the bit.

Tips for Testing Bits

As mentioned previously, the Receive call and three send calls return status information to a 40-byte user buffer. In some cases, you may want to test bits in one-byte binary items. A sample technique for COBOL programs follows:

```
IDENTIFICATION DIVISION.
PROGRAM-ID.
    BITS.
AUTHOR.
    R-FRIED.
INSTALLATION. DIGITAL EQUIPMENT CORPORATION.
DATE-WRITTEN.
DATE-COMPILED.

ENVIRONMENT DIVISION.
CONFIGURATION SECTION.
SOURCE-COMPUTER.      PDP-11.
OBJECT-COMPUTER.      PDP-11.

DATA DIVISION.
WORKING-STORAGE SECTION.
01  BITS-TESTER PIC 9(4) COMP VALUE 0.
01  TEST-BYTE REDEFINES BITS-TESTER PIC X.
01  BIT-WORK PIC S9(4) COMP.
01  BIT-VALUE PIC 9(4) COMP.
01  BIT-SUB PIC 9(4) COMP.
01  THE-BITS.
    03 0-BIT PIC X.
    03 1-BIT PIC X.
    03 2-BIT PIC X.
    03 3-BIT PIC X.
    03 4-BIT PIC X.
    03 5-BIT PIC X.
    03 6-BIT PIC X.
    03 7-BIT PIC X.
01  BIT-ARRAY REDEFINES THE-BITS.
    03 A-BIT PIC X OCCURS 8.
01  BYTE-YOU-WANT-TO-TEST PIC X.

PROCEDURE DIVISION.
TESTER SECTION.
SBEGIN.
    MOVE BYTE-YOU-WANT-TO-TEST TO TEST-BYTE.
    PERFORM COMPUTE-BITS.
```

Then, for example, to test that the two low-order bits are set —

```
IF 0-BIT = "1" AND 1-BIT = "1" ...  
COMPUTE-BITS SECTION.  
STARTIT.  
    MOVE ALL "0" TO THE-BITS.  
    MOVE 128 TO BIT-VALUE.  
    MOVE 8 TO BIT-SUB.  
LOOP.  
    PERFORM COMP-BITS.  
    SUBTRACT 1 FROM BIT-SUB.  
    IF BIT-SUB = 0 OR BIT-TESTER = 0  
        GO TO GET-OUT.  
    GO TO LOOP.  
COMP-BITS.  
    SUBTRACT BIT-VALUE FROM BITS-TESTER  
        GIVING BIT-WORK.  
    IF BIT-WORK NOT NEGATIVE  
        MOVE "1" TO A-BIT (BIT-SUB)  
        MOVE BIT-WORK TO BITS-TESTER.  
    DIVIDE 2 INTO BIT-VALUE.  
GET-OUT.  
EXIT.
```

7.1.2 Compiling and Linking

The DECnet/E subprograms are stored as object modules in the RSTS/E system library (LB:SYSLIB.OLB). References to the subprograms will be resolved automatically during the task building process unless you give specific directions. Follow your normal procedure for compiling, merging ODLs, and task building, as described in the *PDP-11 COBOL User's Guide*.

7.2 Declare Receiver (MDCL) — Privileged (Local and DECnet)

The MDCL call must be executed before a program can receive any messages or send any network messages. (Local messages can be sent without first executing this call.) The call defines an identifying object type code or name (or both) and any restrictions on incoming messages. The monitor associates this information with the RSTS/E job number, setting up a Receiver ID Block for the job, as described in Section 5.2.3. Only one MDCL call can be in effect at a time for a particular job. It remains in effect until a Remove Receiver call (MREM) is executed for the job.

COBOL Call Format:

CALL "MDCL" USING *errval*, *name*, *objtyp*, *access*, *bmax*, *mmax*, *lmax*

Argument Descriptions:

errval

The *errval* argument is a one-word COMP data item. It is set to zero if the MDCL call succeeds. If an error occurs, *errval* is set to one of the error codes listed under "Possible Errors" at the end of this section.

name

The logical name is a data-name referring to the first character position of a 6-byte area containing one to six characters. Valid characters are uppercase alphanumerics and the special characters "\$" (dollar sign), "." (period), and "_" (underline). If less than six characters are used, the name must be left-justified in the 6-byte area and the remaining bytes must be padded with spaces. (The VALUE IS clause or MOVE statement will do this automatically.) Leading or embedded spaces are invalid. If no name is to be specified (that is, the logical name is null), this area should contain six spaces.

The name is used to identify the calling program for queuing messages from local or network programs. Local programs can address the calling program by the logical name. Network programs can use either the logical name or the object type code, if one is declared.

For local or network programs to access the calling program by logical name, the name must be nonnull and unique to the node. That is, only one program at the node can declare its identity with that name at any given time. NSP uses the logical name to identify the calling program for queuing messages from local or network programs. If the logical name is null, only network access by object type code is allowed. (See the following discussion of *logical name*, *objtyp*, and *access*.)

objtyp

This one-word COMP data item specifies the object type code, another form of network addressing. As discussed in Section 5.2, the object type defines a service that the calling program performs. If the program issuing the MDCL call is addressed by object type alone (*name* is null and *access* indicates network access only), multiple copies can execute at the same time in the

RSTS/E environment. Unlike the logical name, the object code need not be unique within the node — multiple copies of a program can declare their identity with the same type code at the same time. If *name* is null, *objtyp* cannot be zero. (Network addressing is discussed in detail in Section 6.1.)

Acceptable values for the object type code range from 0 through 255. The value 0 indicates that other programs in the network will never address the calling program by object type code. The range 1 to 127 is reserved for DECnet use (see Appendix A). The range 128 to 255 is available for definition and use by a network installation.

access

This one-word COMP data item contains the sum of five bit values ($l + p + n + o + s$) that are used to determine access to the declaring program and to modify certain aspects of the message operations.

Three of these bit values (l , p , and n) limit the types of senders that can queue messages for the calling program. They do not, however, limit the messages that the program can send.

If $l = 0$, messages from local senders will not be queued for the calling program. (Local senders who use the network functions are considered network senders in this context.) If $l = 1$, messages from local senders will be queued.

If $p = 0$, incoming messages from both local privileged and nonprivileged senders will be queued for the calling program. If $p = 2$, incoming messages from local senders will be queued only if the sender is privileged. (This bit is ignored if $l = 0$, that is, if no local senders are allowed.)

If $n = 0$, incoming requests for logical links are rejected by NSP. If $n = 4$, incoming Connect Initiate Messages are passed on to the receiver, subject to the declared link maximum ($lmax$). In this case, the program must provide its own protection from unauthorized access, if necessary.

Table 7-1 summarizes access code values and the access permitted for each.

Table 7-1: Types of Access Permitted

| Access Code ($n+p+l$) | Network Logical Links | Local Senders Privileged | Local Senders Allowed |
|----------------------------|--------------------------|-----------------------------|--------------------------|
| 0 | no | no | no |
| 1 | no | no | yes |
| 2 | no | no | no |
| 3 | no | yes | yes |
| 4 | yes | no | no |
| 5 | yes | no | yes |
| 6 | yes | no | no |
| 7 | yes | yes | yes |

Bit value *o* is used to regulate single-link programs, modifying the function of the program's declared link maximum (*lmax*). If *o* = 0, incoming requests for logical links do not affect the link maximum. If *o* = 8, however, after queuing an incoming Connect Initiate Message to the receiver, NSP sets the program's link maximum to zero, inhibiting all further incoming connect requests. This, in effect, modifies the meaning of the link maximum from "one link at a time" to "one link per program execution."

By setting this "one-shot" bit to 8, a program prevents the possibility that NSP will queue a Connect Initiate Message for it after its one logical link is disconnected and before it issues a Remove Receiver call.

The DECnet/E FAL utility, for example, sets the one-shot bit to 8 and declares a link maximum of 1. Once FAL is started as the result of an incoming request from a remote NFT utility, NSP queues the Connect Initiate Message and zeros the link maximum. Thus, there is no possibility that another incoming connect request could be queued for that copy between the time that NFT breaks the first link and the time that FAL issues its Remove Receiver call. This ensures that each incoming connect request starts a new copy of FAL.

Bit value *s* is used to modify the function of the RSTS/E conditional Sleep monitor call. If *s* = 0, any unreceived messages queued for the program will block the execution of a conditional Sleep call. This is the normal RSTS/E function. If *s* = 16, however, the RSTS/E monitor will not check the program's message queue when determining whether or not it should in fact suspend program execution. Several other conditions (such as receiving a delimiter on an open terminal) can still block the sleep, but a pending message will not.

logical name, objtyp, and access

NSP checks the *name*, *objtyp*, and *access* arguments to ensure that they are consistent. For example, a null logical name and an access allowing local senders are inconsistent since local senders can address the program by logical name. NSP checks these arguments as shown in Table 7-2. If the arguments are invalid, NSP returns an error code of 18 in the *errval* argument. If the arguments are valid, NSP allocates a small buffer to hold the information passed. If no buffers are available, the call fails with *errval* = 4. A retry may succeed.

bmax

Until a Receive call (MRCV) is executed, a pending message occupies system buffer space. One 16-word buffer from the monitor's buffer pool is used for user- or DECnet-defined parameters and other system-specific information. Additional buffer space for the message data is usually allocated from the extended buffer pool. If an extended pool does not exist, or no space is available there, the monitor's buffer pool is used.

Because the monitor pool is a critical system resource, the program must set a limit on the amount of space to be used on its behalf. The integer argument *bmax* sets a limit (1 to 32767 bytes) on the total monitor pool space to be used for the user data portion of messages.

Table 7-2: System Validation of Name, Object, and Access Arguments

| Access Code (n+p+l) | Senders Permitted | Object Type | Logical Name |
|------------------------|----------------------|----------------|---|
| 0,2 | None | Ignored | Ignored. |
| 1,3 | Local only | Ignored | Must be nonnull and unique. |
| 4,6 | Network only | Zero | Must be nonnull and unique. |
| | | Nonzero | If the logical name is null (at least one leading zero byte), the only access permitted is by object type code. Multiple copies of the program can coexist. |
| | | | If the logical name is nonnull, it must be unique at the local node. Network senders can refer to the calling program by name or by object type code. Only one copy of the program using this logical name can execute at any given time. |
| 5,7 | Network/Local | Any value | Must be nonnull and unique so that local senders can refer to the calling program by logical name. Network senders can use the logical name or object type code (if <i>objtyp</i> ≠ 0). Only one copy of the program using this logical name can execute at any given time. |

When this buffer maximum is reached, Local Data Messages are no longer queued for the program. NSP returns an error to any local program that tries to send a Local Data Message at this point. Network messages are queued regardless of the declared buffer maximum, but any monitor pool space used is counted against the maximum.

A zero or negative value here indicates that the monitor's buffer pool is not to be used for the user data portion of pending Local Data Messages.

mmax

The message maximum (*mmax*), a one-word COMP data item ranging from 0 to 255, limits the number of messages that will be queued for the calling program.

NSP keeps count of the total number of messages queued for each declared receiver. When one of these counters equals or exceeds the message maximum set by a receiver:

1. Incoming Local Data Messages are no longer queued for the receiving program. An error is returned to the local sender.
2. Incoming Connect Initiate Messages cause NSP to determine whether another copy of the program can be started automatically (see Section

- 6.1.2). If not, NSP rejects the connection and the Connect Initiate Message is not queued. A Connect Reject Message is returned to the sender.
3. An incoming Network Data Message causes NSP to inhibit incoming data messages on that particular logical link and to negatively acknowledge (NAK) the message. This forces the remote system to retransmit when flow is reenabled. (See Section 6.2.2 for further details on backpressure flow control.)

lmax

The link maximum (*lmax*), a one-word COMP data item ranging from 0 to 63, limits the number of incoming requests for logical links for the program. If the number of links currently active is greater than or equal to *lmax* when a remote Connect Initiate Message is received for the program, the local NSP checks to determine whether another copy of the program can be started automatically (see Section 6.1.2). If not, NSP rejects the connection and the Connect Initiate Message is not queued. A Connect Reject Message is returned to the sender.

By declaring a small link maximum, a program can avoid the overhead of responding to remote Connect Initiate Messages when it is known beforehand that the program can only handle a limited number of logical links.

A zero value for *lmax* has the same effect as setting *n* = 0 in the *access* argument. That is, NSP rejects all incoming requests for logical links to the calling program. The *lmax* argument does not, however, limit the number of logical links initiated by the program. That is, it does not stop the program from sending Connect Initiate Messages.

Possible Errors:

| errval | Error Text and Meaning |
|---------------|--|
| 1 | BAD DIRECTORY ON DEVICE Network operations have been terminated due to an internal error. Notify the system manager. |
| 3 | ACCOUNT OR DEVICE IN USE The calling job already exists in the system's list of declared receivers. This error can indicate a logic error in the program or that a previous program running under the same job number failed to remove itself from the receiver list before terminating. In the latter case, a Remove Receiver call should be issued, followed by another Declare Receiver call. (It is common practice to issue a Remove Receiver call immediately before the Declare Receiver call.) |
| 4 | NO ROOM FOR USER ON DEVICE There were no small buffers available to hold the arguments passed with the Declare Receiver call. Since the system's use of small buffers is dynamic, a retry may succeed. |

(continued on next page)

errval Error Text and Meaning

10 PROTECTION VIOLATION

The calling program must be privileged at the time the Declare Receiver call is issued.

16 NAME OR ACCOUNT NOW EXISTS

The logical name passed is being used by another job.

18 ILLEGAL SYS() USAGE

The parameters passed are inconsistent. See *logical name*, *object*, and *access* discussion.

Example:

The following example shows a Declare Receiver call allowing both network and local access. Both a logical name (CNET) and an object type code (142) are specified. The buffer maximum is set to 512 and the message maximum is set to 5. No more than 512 bytes of user message data from incoming Local Data Messages will be stored in monitor pool space. And, if more than 5 messages accumulate for the program, an incoming request for a connection will cause NSP to check to determine whether another copy of the program can be started automatically. If not, NSP will reject the link and any incoming data messages will cause NSP to inhibit incoming message flow on all the program's logical links because of the backpressure condition. The link maximum is set to 2. When the program has established two logical links, an incoming Connect Initiate Message will cause NSP to check to determine whether another copy of the program can be started. If not, NSP will reject the link.

(Data Division)

```
01        STAT PIC 9(4) COMP.
01        DECLARE-ITEMS.
          02        NET-NAME PIC X(6) DISPLAY VALUE IS "CNET ".
          02        OBJTYP PIC 9(4) COMP VALUE IS 142.
          02        ACCESSM PIC 9(4) COMP VALUE IS 5.
          02        BMAX PIC 9(4) COMP VALUE IS 512.
          02        MMAX PIC 9(4) COMP VALUE IS 5.
          02        LMAX PIC 9(4) COMP VALUE IS 2.
          .
          .
          .
```

(Procedure Division)

```
CALL "MDCL" USING STAT, NET-NAME, OBJTYP, ACCESSM,
              BMAX, MMAX, LMAX.
```

7.3 Remove Receiver (MREM)

(Local and DECnet)

The MREM call removes the specified job from the system's list of declared receivers. All pending messages are discarded and any active logical links are aborted. This call should be executed at the completion of message processing. This prevents unwanted messages from accumulating in the queue of pending messages.

Privileged programs can remove other RSTS/E jobs with this call, although normally only system utility programs would do this.

COBOL Call Format:

CALL "MREM" USING *errval*, *jobx2*

Argument Descriptions:

errval

The *errval* argument is a one-word COMP data item. It is set to zero if the MREM call succeeds. If an error occurs, *errval* is set to one of the error codes listed under "Possible Errors" at the end of this section.

jobx2

This argument is a one-word COMP data item. A zero value for *jobx2* removes the calling job. Privileged programs can remove another local RSTS/E job by specifying job number times two for this argument.

Possible Errors:

| <i>errval</i> | Error Text and Meaning |
|---------------|--|
| 1 | BAD DIRECTORY ON DEVICE Network operations have been terminated due to an internal error. Notify the system manager. |
| 10 | PROTECTION VIOLATION The caller is nonprivileged and has attempted to remove another job (that is, <i>jobx2</i> is nonzero). |
| 18 | ILLEGAL SYS() USAGE The <i>jobx2</i> argument was odd. <i>Jobx2</i> must be zero to remove the calling program or job number times two to remove another job. |

Example:

The following example shows a typical Remove Receiver call, removing the calling job from the system's list of declared receivers.

(Data Division)

```
01      ERRVAL PIC 9(4) COMP.  
01      JOBX2 PIC 9(4) COMP VALUE IS 0.  
      .  
      .  
      .
```

(Procedure Division)

```
      CALL "MREM" USING ERRVAL , JOBX2.
```

7.4 Get Local Node Parameters (NTLN)

(DECnet)

The NTLN call returns information to the calling program concerning the local node's network parameters. Data returned includes the node's name, address and alias (if any), and the default user's account to be used when an incoming Connect Initiate Message does not specify accounting information. (See the *DECnet/E System Manager's Guide* for a discussion of alias and Section 6.1.4 of this manual for a discussion of the default user's account.)

COBOL Call Format:

CALL "NTLN" USING *errval*, *buffer*, *buflen*, *bufoff*

Argument Descriptions:

errval

The *errval* argument is a one-word COMP data item. It is set to zero if the NTLN call succeeds. If an error occurs, *errval* is set to one of the error codes listed under "Possible Errors" at the end of this section.

buffer

This argument is a data-name referring to the first character position (byte) of the buffer to receive the local node information. The data-name can refer to a single data item or an aggregate of data items (for example, a level 01 group name). Any and all PICTURE categories can be used in the fields referenced by the data-name. The returned data can be offset from the beginning of the buffer (see *bufoff*).

When the NTLN call completes successfully, data is returned to the buffer in the following format:

| Bytes | Contents |
|-------|--|
| 0-5 | Local node name, consisting of 1 to 6 uppercase, alphanumeric ASCII characters. The name contains at least one alphabetic character. A name shorter than 6 characters is left-justified and padded to 6 characters with spaces. |
| 6-7 | Local node address, ranging from 1 to 255. |
| 8-13 | Local node alias, consisting of 1 to 6 uppercase, alphanumeric ASCII characters. The alias contains at least one alphabetic character. An alias shorter than 6 characters is left-justified and padded to 6 characters with spaces. If no alias is defined, a field of 6 spaces is returned. |
| 14-15 | Local default project-programmer number. Byte 14 = programmer number; byte 15 = project number. If no default account is defined, both bytes are zero. |

This one-word COMP data item defines the number of bytes to be returned to the buffer specified with *buffer*. If *buflen* is less than 16 bytes, the call returns only as much data as there is room for. For example, if *buflen* = 8 bytes, the call returns only the local node name and address. Sixteen bytes are required to receive all the local node parameters.

This one-word COMP data item defines an offset, in bytes, from the beginning of the buffer. For example, if offset = 20, the returned data begins in byte 21 of the buffer.



The extended COBOL interface makes no checks for inconsistencies in the *buffer*, *buflen*, and *bufoff* values specified. A program could, for example, specify a *buflen* of 16 and a *bufoff* of 10. If the actual buffer, beginning at *buffer* was only 16 bytes long, the returned data would overwrite 10 bytes beyond the end of the buffer.

| errval | Error Text and Meaning |
|---------------|-------------------------------|
|---------------|-------------------------------|

| | |
|----|--|
| 1 | BAD DIRECTORY ON DEVICE |
| | Network operations have been terminated due to an internal error. Notify the system manager. |
| 62 | NO RUN-TIME SYSTEM |
| | NSP has not been enabled. Normal DECnet calls cannot be executed until the system manager enables NSP. |
| 66 | MISSING SPECIAL FEATURE |
| | DECnet was not installed at system generation time. The network functions cannot be executed. |

Example:

The following example shows a NTLN call used to obtain all the local node parameters. The node name is to be stored in the data item NODE-NAME; the address, in the COMP data item NODE-ADDRESS; the node alias, in the data item NODE-ALIAS; and the default project-programmer number, as a unit in the COMP data item PPN and separately in the COMP data items PROJECT and PROGRAMMER.

(Data Division)

```
01      RETURN-DATA.  
03      NODE-NAME PIC X(6).  
03      NODE-ADDRESS PIC 9(4) COMP.  
03      NODE-ALIAS PIC X(6).  
03      PPN PIC 9(4) COMP.  
03      PROJ-PROG REDEFINES PPN.  
05      PROG-BYTE PIC X.  
05      PROJ-BYTE PIC X.  
01      BUFFER-LENGTH PIC 9(4) COMP VALUE IS 16.  
01      BUFFER-OFFSET PIC 9(4) COMP VALUE IS 0.  
01      ERRVAL PIC 9(4) COMP.  
01      PROGRAMMER PIC 9(4) COMP.  
01      PROJECT PIC 9(4) COMP.  
01      WORK-WORD PIC 9(4) COMP.  
01      WW REDEFINES WORK-WORD.  
03      LOW-BYTE PIC X.  
03      HIGH-BYTE PIC X.  
      .  
      .  
      .
```

(Procedure Division)

```
      CALL "NTLN" USING ERRVAL, RETURN-DATA, BUFFER-LENGTH,  
      BUFFER-OFFSET.  
      MOVE 0 TO WORK-WORD.  
      MOVE PROG-BYTE TO LOW-BYTE.  
      MOVE WORK-WORD TO PROGRAMMER.  
      MOVE PROJ-BYTE TO LOW-BYTE.  
      MOVE WORK-WORD TO PROJECT.
```

7.5 Log User Event (NTEV) — Privileged

(DECnet)

The NTEV call permits a user-written program to queue an event to the system event processor for logging. Before the event is logged, it is time-stamped by the system, in the standard Network Information and Control Exchange (NICE) protocol format. The system manager can use the normal network management SET and CLEAR LOGGING commands to control the filtering of these events. Optional data can accompany a user event.

NOTE

This system call performs a highly specialized function which requires a great deal of special knowledge of DECnet and the DIGITAL Network Architecture (DNA). It is provided as a convenience for the sophisticated network user and is not intended for normal network programs.

COBOL Call Format:

```
CALL "NTEV" USING errval,evmod,evcls,evtyp,enval,  
                  buffer,buflen,bufoff
```

Argument Descriptions:

errval

The *errval* argument is a one-word COMP data item. It is set to zero if the NTEV call succeeds. If an error occurs, *errval* is set to one of the error codes listed under "Possible Errors" at the end of this section.

evmod

This one-word COMP data item specifies what the system should do if it is unable to log the user event due to lack of resources. If *evmod* is set to zero, the call terminates with a status return of NOROOM (*errval* = 4); the program can retry logging the event after a short delay. If *evmod* = 1, the system logs a "missed event" event message, and the call terminates with a successful status (*errval* = 0).

Note that in the second case, the user data is *not* logged. The system merely logs the fact that it has lost some unidentified data. This is identical to the procedure followed with DECnet internal events when the event queue is full or there are no more small buffers.

evcls

This one-word COMP data item specifies the class of the user event to be logged. Events are divided into classes, according to the DNA layer from which they originate. Classes 480 to 511 are reserved for customer-specific events. However, DECnet/E supports only class 480. Thus, *evcls* must be 480.

evtyp

This one-word COMP data item specifies the type of event within the class specified by *evcls*. The event type must be in the range of 0 to 31.

entyp

This one-word COMP data item specifies the entity type with which the event is concerned. There are four entity types supported by DECnet/E, defined as follows:

- 1 No specific entity
- 0 Node
- 1 Line
- 3 Circuit

enval

This one-word COMP data item associates a value to the entity specified by the *entyp* argument. If the entity is a node, *enval* should be set to the node address. If the entity is a line or a circuit, *enval* should be set to the address of the Device Data Block (DDB) maintained for the device.

buffer

This argument is a data-name referring to the first character position of the buffer containing the optional user data to be processed with the event. The data-name can refer to a single data item or an aggregate of data items (for example, a level 01 group name). Any and all PICTURE categories can be used in the fields referenced by data-name. The data to be processed can be offset from the beginning of the buffer (see *bufoff*).

If this data is included, it must be in NICE protocol format. Basically, this means that the data should be formatted in a 3-part field: a parameter type, a data type, and a parameter value. DECnet parameter types can be used if the event concerns a node, a circuit, or a line. In this case, the event logger will output standard descriptive text for the specific parameter. If a DECnet parameter type is not used, parameter types in the range of 3900 to 4095 can be specified. (See the *DIGITAL Network Architecture, Network Management Functional Specification*, Version 3.0.0, for further information on event definitions and parameters.)

Up to 200 bytes of data can be passed to the event logger for processing.

buflen

This one-word COMP data item defines the number of bytes of data in the buffer specified with *buffer*. Up to 200 bytes of data can be passed. If no data is specified, *buflen* must be 0.

This one-word COMP data item defines an offset, in bytes, from the beginning of the buffer. For example, if *bufoff* = 20, the data to be processed begins in byte 21 of the buffer.



Possible Errors:

Extended COBOL Interface Call Formats 7-19

Example:

The following example illustrates the logging of a user event, class 480, and type 3. The entity is node number 135. The optional data passed for processing consists of an *lla* parameter number (2130) and value. A "missed event" is logged if the event itself cannot be. Note that even though 6 bytes are allocated for the parameter value, only 2 are used.

```
01  MODIFIER PIC 9(4) COMP VALUE IS 1.
01  EVENT-CLASS PIC 9(4) COMP.
01  EVENT-TYPE PIC 9(4) COMP.
01  ENTITY-TYPE PIC 9(4) COMP.
01  ENTITY-VALUE PIC 9(4) COMP.
01  PARAMETER-DATA.
    03  FILLER PIC X(6).
    03  PARAMETER-NUMBER PIC 9(4) COMP.
    03  DATA-LENGTH PIC X.
    03  DATA-VALUE-1 PIC 9(4) COMP.
    03  DATA-VALUE-2 PIC 9(4) COMP.
    03  DATA-VALUE-3 PIC 9(4) COMP.
01  BUFFER-LENGTH PIC 9(4) COMP.
01  BUFFER-OFFSET PIC 9(4) COMP VALUE IS 6.

01  WORK-WORD PIC 9(4) COMP.
01  WW REDEFINES WORK-WORD.
    03  LOW-BYTE PIC X.
    03  HIGH-BYTE PIC X.
    .
    .
    .
```

(Procedure Division)

```
MOVE 480 TO EVENT-CLASS.
MOVE 3 TO EVENT-TYPE.
MOVE 0 TO ENTITY-TYPE.
MOVE 135 TO ENTITY-VALUE.
MOVE 2130 TO PARAMETER-NUMBER.
MOVE 2 TO WORK-WORD.
MOVE LOW-BYTE TO DATA-LENGTH.
MOVE 5 TO BUFFER-LENGTH.

CALL "NTEV" USING ERRVAL, MODIFIER, EVENT-CLASS,
EVENT-TYPE, ENTITY-TYPE, ENTITY-VALUE,
PARAMETER-DATA, BUFFER-LENGTH, BUFFER-OFFSET.
```

The NTEV call in the previous example generates the following event message:

```
Event type 480.3
Occurred 14-Dec-81 15:31:39.7 on node 135 (GROK)
Node address 135 (GROK)
Parameter #2130 = 40960
```

7.6 Send Calls

There are nine separate send calls — one local and eight network. Each of these calls causes NSP to format a message that is then either queued to a local program or transmitted across the network to a remote program.

7.6.1 Send Local Data Message (MSLD)

(Local)

The MSLD call sends up to 532 bytes of user data to another program on the local RSTS/E system. The other program must be a declared receiver of local messages. If the other program has declared that only privileged local senders are allowed, the calling program must be privileged or the MSLD call will fail with an error. The receiving program is identified in the MSLD call either by its job number or by logical name with which it declared itself a receiver.

COBOL Call Format:

CALL "MSLD" USING *errval*, *jobx2*, *name*, *buffer*, *buflen*, *bufoff*,
parptr, *parlen*

Argument Descriptions:

errval

The *errval* argument is a one-word COMP data item. It is set to zero if the MSLD call succeeds. If an error occurs, *errval* is set to one of the error codes listed under "Possible Errors" at the end of this section.

jobx2

The *jobx2* argument is a one-word COMP data item. If *jobx2* is zero, the message data is sent to the local program identified by the *name* argument. Otherwise, the message data is sent to the program whose job number (times two) is equal to the value of *jobx2*. For example, if *jobx2* = 8, the message is sent to job 4.

name

When the *jobx2* argument = 0, this argument is a data-name referring to the first character of a 6-character field containing the logical name identifying the intended local receiver. The receiver must have declared its identity with this name in its Declare Receiver call. It is a one- to six-character string. Valid characters are uppercase, alphanumerics and the special characters "\$" (dollar sign), "." (period), and "_" (underline). The characters are left-justified in the string and (if necessary) padded to six characters with spaces. (The VALUE IS clause or MOVE statement does this automatically.) Leading or embedded spaces are invalid. This argument is ignored when *jobx2* ≠ 0.

buffer

This argument is a data-name referring to the first character position (byte) of a buffer containing the user message data to be sent. The data-name can refer to a single data item or an aggregate of data items (for example, a level 01 group name). Any and all PICTURE categories can be used in the fields referenced by data-name. The message data can be offset from the beginning of the buffer (see *bufoff*).

This one-word COMP data item defines the number of bytes to be sent from the buffer. Possible values for *buflen* range from 0 through 512.

This one-word COMP data item defines an offset, in bytes, from the beginning of the buffer. For example, if offset = 20, the user data begins in byte 21 of the buffer.



parptr

parlen

This one-word COMP data item defines the number of bytes to be sent from the parameter area defined by *parptr*. The *parlen* value can range from 0 through 20.

| errval | Error Text and Meaning |
|---------------|-------------------------------|
|---------------|-------------------------------|

(continued on next page)

errval Error Text and Meaning

10 PROTECTION VIOLATION

Some access violation has occurred. Either the receiver does not allow any local senders or the sender is nonprivileged and the receiver allows only privileged senders.

18 ILLEGAL SYS() USAGE

The value of *jobx2* is odd. It must be 0 or the receiver's job number times two.

31 ILLEGAL BYTE COUNT FOR I/O

Either the *buflen* value or the *parlen* value is invalid. The *buflen* argument can range from 0 through 512. The *parlen* argument can range from 0 through 20.

32 NO BUFFER SPACE AVAILABLE

System buffers are currently not available to store this message for the intended local receiver. A later retry may proceed without error.

Example:

The following example shows a local send to a program that has issued a Declare Receiver call with the logical name LOCAL. A total of 512 bytes are sent from BUFFER. No parameter message data is sent.

(Data Division)

```
01        BUFFER PIC X (512).
01        STAT PIC 9(4) COMP.
01        SEND-LOCAL-ITEMS.
03           REM-NAME PIC X(6) VALUE IS "LOCAL ".
03           JOBX2 PIC 9(4) COMP VALUE IS 0.
03           BUFLN PIC 9(4) COMP VALUE IS 512.
03           PARLEN PIC 9(4) COMP VALUE IS 0.
         .
         .
         .
```

(Procedure Division)

```
CALL "MSLD" USING STAT, JOBX2, REM-NAME, BUFFER,
         BUFLN, BUFFER, PARLEN.
```

7.6.2 Send Connect Initiate Message (NTCI)

(DECnet)

The NTCI call requests a logical link to another program. The call establishes the user link address that the calling program will use to refer to the link. The other program is identified with a Connect Data Block. Other arguments specify the flow control option and maximum amount of user data for received Network Data Messages. Up to 16 bytes of user data can be sent with the NTCI call.

COBOL Call Format:

CALL "NTCI" USING *errval*, *ula*, *llmod*, *rmax*, *buffer*, *buflen*, *bufoff*

Argument Descriptions:

errval

The *errval* argument is a one-word COMP data item. It is set to zero if the NTCI call succeeds. If an error occurs, *errval* is set to one of the error codes listed under "Possible Errors" at the end of this section.

ula

The user link address (*ula*) is a one-word COMP data item in the range of 1 to 255. Later calls to send and receive messages use this value to refer to the logical link (assuming that the link being requested by the NTCI call is accepted by the remote NSP and the remote program).

The *ula* must be unique within the program at any given time. That is, one program cannot have two different logical links with the same *ula* at the same time.

llmod

The one-word COMP data item *llmod* indicates the type of flow control requested by the calling program to control incoming data from the remote program. Acceptable values and meanings are:

- 0 No flow control
- 1 Segment flow control
- 2 Message flow control

Flow control is described in detail in Chapter 6. Remember that each program on a logical link selects its own flow control. The *llmod* value is the selection made by the program issuing the call. If *llmod* is nonzero, the program must issue Link Service (NTLS) calls to request Network Data Messages from the other program.

rmax

A program that has limited buffer space but is not designed to process large messages in small pieces, can impose a limit on the amount of user message data it is willing to receive on a logical link. The *rmax* argument, a one-word COMP data item, specifies this limit in bytes.

The size of the receive buffers allocated by NSP will itself limit the amount of data that can be passed in a single Network Data Message. If the maximum size set for NSP by Network Management is not large enough to handle the receive maximum as specified in *rmax*, the local NSP will alter the maximum to the smaller limit before forwarding the Connect Initiate Message to the remote node. If *rmax* = 0, NSP will supply the size of its receive buffers as a default value. (The calling program will be informed of any such change by a field in the Connect Confirm Message from the target program when the link is accepted.)

The remote NSP uses this value to establish a transmit maximum for its end of the logical link. If the remote system is also a RSTS/E system, this transmit maximum is passed on to the remote program. The remote program must limit the amount of user message data it sends over the logical link according to this value.

See the discussion of the *buflen* argument for the Receive call (MRCV) in Section 7.7 for a description of how to process large messages in small pieces.

buffer

This argument is a data-name referring to the first character position (byte) of a buffer containing the Connect Data Block and optional user message data. The 120-byte Connect Data Block and up to 16 bytes of message data are assumed to be in the buffer in contiguous bytes. They can be offset from the beginning of the buffer (see *bufoff*).

The Connect Data Block identifies the target program for the requested logical link. Figure 7-1 shows the format of the Connect Data Block. The fields within are described in Table 7-3.


Up to 16 bytes of user data can follow the Connect Data Block. Any PICTURE categories can be used.

buflen

This one-word COMP data item defines the number of bytes to be sent. Since the Connect Data Block is 120 bytes long and the optional user data can be from 0 to 16 bytes long, acceptable values for *buflen* range from 120 through 136.

bufoff

This one-word COMP data item defines an offset, in bytes, from the start of the buffer from which the data is to be sent. For example, if *bufoff* = 20, the Connect Data Block begins in byte 21 of the buffer.



NOTE

The extended COBOL interface makes no checks for inconsistencies in the *buffer*, *buflen*, and *bufoff* values specified. If, for example, a program's buffer is 100 bytes long, with an offset of 20 and a length of 136, 136 bytes will be sent: the last 80 bytes of the buffer, and 56 bytes of whatever is beyond the buffer in memory.

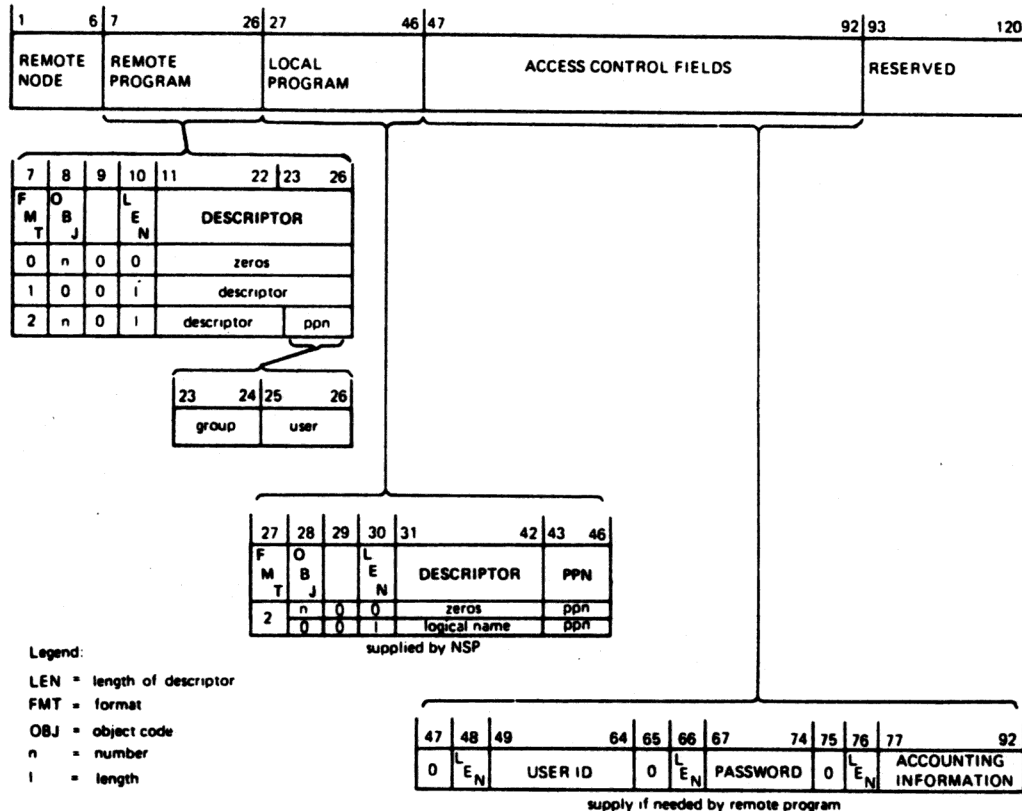


Figure 7-1: Format of Connect Data Block on a Send Connect Initiate Message

Table 7-3: Format of Connect Data Block on a Send Connect Initiate Message

| Bytes | Content |
|-------|--|
| 1-6 | Remote Node. One to six uppercase, alphanumeric ASCII characters naming the remote node to which the Connect Initiate Message is to be sent. The name must contain at least one alphabetic character, and names shorter than six characters must be left-justified and padded to six characters with spaces. If no node name is to be specified (that is, the request is directed to a program on the local node), these six bytes should be set to spaces. |

(continued on next page)

Table 7-3 (Cont.): Format of Connect Data Block on a Send Connect Initiate Message

| | |
|--------------|--|
| Bytes | Content |
| 7-26 | Remote Program. Identifies the remote program to which the Connect Initiate Message is directed as either an object type (general function) or a specific entity (program, task, or process). This identification is given in one of three formats: |

Format 0

The remote program is identified by object type code alone. The fields are:

| | |
|--------------|--|
| Bytes | Content |
| 7 | One-byte binary value of zero, to indicate a format 0 name. |
| 8 | One-byte binary value giving the object type code of the remote program. |
| 9-26 | Zero (Not used). |

Format 1

The remote program is identified by a descriptor. The fields are:

| | |
|--------------|--|
| Bytes | Content |
| 7 | One-byte binary value of one, to indicate a format 1 name. |
| 8 | Normally zero (one-byte binary value). This byte should be zero unless the target DECnet system allows identification by both name and object type code on incoming connect requests (DECnet/E does not). If nonzero, this byte is interpreted as the object type code of the remote program for which the Connect Initiate Message is intended. |
| 9 | Reserved - must be one-byte binary zero. |
| 10 | One-byte binary value giving the length of the descriptor, 0 to 16 bytes. Gives the number of characters to be interpreted as the program descriptor in the following field. |
| 11-26 | Remote program descriptor (left-justified). For a DECnet/E program, this is the logical name that the remote program defined in its Declare Receiver call. For other systems, it is the descriptor used to register the program for network operations. |

Format 2

This format allows identification of the remote program by name or object type code and by the group and user codes under which the remote program is accessible on the remote system. (The group and user codes refer to what is called the project-programmer number on some DIGITAL systems, including RSTS/E. Other systems refer to these codes as the user identification code, or UIC.) DECnet/E systems use only the object type code or descriptor to find or start the target program for which a Connect Initiate Message is intended. The group and user codes are simply passed on to the target program for its own checking. Other DECnet systems may handle an incoming connect request in format 2 differently.

The subfields are defined as follows:

| | |
|--------------|--|
| Bytes | Content |
| 7 | One-byte binary value of two, to indicate a format 2 name. |
| 8 | One-byte binary object type code of the remote program, if applicable. If zero, the descriptor must be used. |

(continued on next page)

Table 7-3 (Cont.): Format of Connect Data Block on a Send Connect Initiate Message

| Bytes | Content |
|-------|--|
| 9 | Reserved - must be one-byte binary zero. |
| 10 | Length of descriptor, 0 to 12 bytes. Binary value giving the number of characters to be interpreted as the program descriptor in the following field (bytes 11-22). |
| 11-22 | Remote program descriptor, if this is used instead of object type code. The descriptor should be left-justified in the field. |
| 23-24 | Remote group code. One-word binary value giving the group number under which the remote program is running or is to be started. Corresponds to the project number in a RSTS/E project-programmer number. This value is not used by DECnet/E systems but is simply passed on to the receiving program. |
| 25-26 | Remote user code. One-word binary value giving the user number under which the remote program is running or is to be started. Corresponds to the programmer number in a RSTS/E project-programmer number. This value is not used by DECnet/E systems but is simply passed on to the receiving program. |
| 27-46 | <p>Local Program. This field is filled in by the local NSP using information that the local program specified in its Declare Receiver call. Any values passed by the user program in this area are ignored.</p> <p>The information is passed to the remote system in format 2, as described previously. Byte 27 is filled in with a one-byte binary value of 2. Byte 29 is always zero. If the program declared its identity with object type code alone, the object type code is passed in byte 28 (one-byte binary value) and bytes 30-42 are zero. If the program declared its identity with a logical name, or with a logical name and object type code, the local name is supplied in bytes 31-42 and byte 28 is a one-byte binary value of zero. The length of the logical name is given in byte 30. The project-programmer number under which the program is running is passed in bytes 43-46 as the group and user codes. The project number is in byte 43. The programmer number in byte 45. Both are in one-byte binary format.</p> |
| 47-92 | <p>Access Control Fields. These bytes can be used to define the calling program's access rights to the remote system or to the remote system's services. The idea is analogous to logging in on the remote system. The most rigorous checking done in a normal log-in on DIGITAL systems involve user identification, password, and account number, so three fields are defined here for those items. However, overall DECnet design does not specifically require use of these fields at all, nor does it require that the fields be used as described. A DECnet/E system simply passes this information on to the receiving program. To determine what, if anything, a non-RSTS/E DECnet system requires in these fields, see the DECnet manual for that system.</p> |

NOTE

The DECnet/E utilities NFT and NETCPY, which use logical links to access remote files and devices, use these fields to establish the local terminal user's right to access remote files and devices. Both utilities prompt the local user for log-in information to be used at the remote system and pass that information along in the Connect Data Block of the Connect Initiate Message. The remote FAL or receiving NETCPY checks this information before allowing the operation requested.

(continued on next page)

Table 7-3 (Cont.): Format of Connect Data Block on a Send Connect Initiate Message

| Bytes | Content |
|--------|---|
| 47 | Reserved - must be zero. |
| 48 | One-byte binary value giving the length of the user identification information specified in the next field. 0 to 16 bytes. |
| 49-64 | User Identification. This field identifies the person or program requesting the logical link. In RSTS/E terms, the user ID is analogous to the project-programmer number used at log-in. This information need only be specified here if it is required by the remote system or program. If it is not required, the length (byte 48) should be set to zero. |
| 65 | Reserved - must be zero. |
| 66 | One-byte binary value giving the length of the password information specified in the next field. 0 to 8 bytes. |
| 67-74 | Password. A password is often used to verify access to a system under a particular user identification. In RSTS/E, this password is analogous to the password used in conjunction with the project-programmer number at log-in. A password supplied here must be acceptable to the remote system or program, if one is required. If not required, the length (byte 66) should be set to zero. |
| 75 | Reserved - must be zero. |
| 76 | One-byte binary value giving the length of the accounting information specified in the following field. 0 to 16 bytes. |
| 77-92 | Accounting information. Many systems require a billing account number in addition to user identification and password. The accounting field is provided for this purpose. Once again, such information need only be supplied here if it is required by the remote system or program. If it is not required, the length (byte 76) should be set to zero. |
| 93-120 | Reserved for future use. Currently ignored by DECnet/E but should be passed as zeros. |

Possible Errors:

| errval | Error Text and Meaning |
|---------------|--|
| 1 | BAD DIRECTORY ON DEVICE Network operations have been terminated due to an internal error. Notify the system manager. |
| 3 | ACCOUNT OR DEVICE IN USE The calling program already has a logical link with the same user link address (<i>ula</i>) as specified in this Connect Initiate Message. A different <i>ula</i> must be used or the existing link must be disconnected. |
| 6 | NOT A VALID DEVICE The node named in the Connect Data Block is not known to the system. |
| 14 | DEVICE HUNG OR WRITE LOCKED The node named in the Connect Data Block is known to the system but is currently inactive. There is no physical communication path to the node. |

(continued on next page)

- 17 **TOO MANY OPEN FILES ON UNIT**
- The total number of logical links allowed for the local node has been reached. A later try may succeed. (The system manager can set a limit on the total number of links. If this limit is inadequate, see the system manager.)
- 18 **ILLEGAL SYS() USAGE**
- One of several possibilities:
1. The calling program is not a declared receiver. A Declare Receiver call must be issued before this call is given.
 2. The Declare Receiver call for the calling program had a null logical name and a zero object type code. There is no way to identify the program to the remote system. Issue a Remove Receiver call and reissue the Declare Receiver call.
 3. An invalid value was detected in one of the following fields.
 1. Remote program field in the Connect Data Block.
 2. One of the three access control fields in the Connect Data Block.
 3. The user link address (*ula*) given was zero. Valid values range from 1 to 255.
 4. The local link modifier (*llmod*) value must be 0, 1, or 2.
 5. The remote descriptor length (byte 10 in the Connect Data Block) was nonzero but the descriptor field was null (all zeros).
- 22 **DISK PACK IS LOCKED OUT**
- NSP cannot create a new logical link because the physical line to the remote node or the local node itself has been scheduled for shutdown by the system manager.
- 31 **ILLEGAL BYTE COUNT FOR I/O**
- The length field passed in *buflen* is invalid. The length of the buffer must be between 120 bytes and 136 bytes for a Connect Initiate Message.
- 32 **NO BUFFER SPACE AVAILABLE**
- System buffers are not currently available to store this message. A later try may succeed.
- 62 **NO RUN-TIME SYSTEM**
- NSP has not been enabled. Normal DECnet calls cannot be executed until the system manager enables NSP.
- 66 **MISSING SPECIAL FEATURE**
- DECnet was not installed at system generation time. The network functions cannot be executed.

Example:

The following example shows a Connect Initiate Message directed to a program named NETRCV at node MIAMI. The Connect Data Block uses a format 1 name to address the remote program. NSP supplies the local program identification and no access control fields are used, so the rest of the Connect Data Block is filled with spaces. A user link address of 27 is established and the receive maximum is set to 512. LLMOD is set to indicate that the local program requests message flow control. Message data is sent following the Connect Data Block.

(Data Division)

```

01  NULLNUM PIC 9(4) COMP VALUE IS 0.
01  NULL REDEFINES NULLNUM PIC X.
01  CONNECT-BUF.
    02  CDB.
        03  REM-NODE PIC X(6).
        03  REM-PROG.
            05  FMT PIC 9(4) COMP VALUE IS 1.
            05  LEN PIC 9(4) COMP.
            05  LEN-REV REDEFINES LEN.
                07  HIGH-BYTE PIC X.
                07  LOW-BYTE PIC X.
            05  DESCRIPTOR PIC X(16) VALUE IS "NETRCV".
        03  LOC-PROG PIC X(20).
        03  ACCESS-FIELDS PIC X(46).
        03  RSRV PIC X(28).
        03  MESSAGE PIC X(7) VALUE IS "OPTION2".
01  TEMP PIC X.
01  STAT PIC 9(4) COMP.
01  ULA PIC 9(4) COMP VALUE IS 27.
01  RMAX PIC 9(4) COMP VALUE IS 512.
01  LLMOD PIC 9(4) COMP VALUE IS 2.
01  BUFLN PIC 9(4) COMP VALUE IS 127.
01  BUFOFF PIC 9(4) COMP VALUE IS 0.

```

(Procedure Division)

```

        MOVE "MIAMI" TO REM-NODE.
        MOVE 6 TO LEN.
        MOVE HIGH-BYTE TO TEMP.
        MOVE LOW-BYTE TO HIGH-BYTE.
        MOVE TEMP TO LOW-BYTE.
        CALL "NTCI" USING STAT, ULA, LLMOD, RMAX,
CONNECT-BUF, BUFLN, BUFOFF.

```

7.6.3 Send Connect Confirm Message (NTCC)

(DECnet)

The NTCC call accepts a remote program's request for a logical link. The call establishes the user link address that the calling program will use to refer to the link. The local link address assigned by the local NSP identifies the particular link being accepted. The flow control option and the receive maximum are specified for the link. Up to 16 bytes of message data can be passed to the remote program.

COBOL Call Format:

CALL "NTCC" USING *errval*, *ula*, *lla*, *llmod*, *rmax*, *buffer*,
buflen, *bufoff*

Argument Descriptions:

errval

The *errval* argument is a one-word COMP data item. It is set to zero if the NTCC call succeeds. If an error occurs, *errval* is set to one of the error codes listed under "Possible Errors" at the end of this section.

ula

The user link address (*ula*) is a one-word COMP data item within the range 1 to 255. Later calls to send and receive messages use this value to refer to the logical link.

The *ula* must be unique within the program at any given time. That is, one program cannot have two different logical links with the same *ula* at the same time.

lla

The local link address (*lla*) is a one-word COMP data item that identifies the link being accepted. It is a number assigned by the local NSP when it received the Connect Initiate Message requesting the link. The local link address is passed to the local program in bytes 5-6 of the received Connect Initiate Message. (Section 7.7.2 describes the information returned for a received Connect Initiate Message.)

llmod

The one-word COMP data item *llmod* indicates the type of flow control requested for this end of the logical link. Acceptable values and meanings are:

- 0 No flow control
- 1 Segment flow control
- 2 Message flow control

Flow control is described in detail in Chapter 6. Remember that both programs on a logical link select their own flow control independent of each other. The *llmod* value is the selection made by the calling program. If *llmod* is nonzero, the program must issue Link Service (NTLS) calls to request Network Data Messages from the other program.

Possible Errors:

| errval | Error Text and Meaning |
|---------------|--|
| 1 | BAD DIRECTORY ON DEVICE Network operations have been terminated due to an internal error. Notify the system manager. |
| 3 | ACCOUNT OR DEVICE IN USE A logical link is currently active for the calling program with the same user link address (<i>ula</i>) as specified in this Connect Confirm Message. A different <i>ula</i> must be used or the existing link must be disconnected. |
| 5 | CAN'T FIND FILE OR ACCOUNT The local link address (<i>lla</i>) does not correspond to any known logical link for the calling program. Either the <i>lla</i> is incorrect or the originating connect request has been cancelled. In the latter case, a Link Abort Message has already been queued for this link. The <i>lla</i> field in the received Link Abort Message identifies the link, since no <i>ula</i> was ever established. |
| 10 | PROTECTION VIOLATION Some procedural error has occurred. Sending a Connect Confirm Message for a link that is not waiting for confirmation will cause this error. For example, a second Connect Confirm Message for the same logical link will cause this error because the link is already established. |
| 18 | ILLEGAL SYS() USAGE One of two possibilities: <ol style="list-style-type: none">1. The calling program is not a declared receiver. A Declare Receiver call must be issued before this call is given.2. An invalid value was detected in one of the following fields.<ol style="list-style-type: none">1. The user link address (<i>ula</i>) given was zero. Valid values range from 1 to 255.2. The local link modifier (<i>llmod</i>) value must be 0, 1, or 2. |
| 31 | ILLEGAL BYTE COUNT FOR I/O The length field passed in <i>buflen</i> is invalid; it must be between 0 and 16 (bytes) for a Connect Confirm Message. |
| 32 | NO BUFFER SPACE AVAILABLE System buffers are not currently available to store this message. A later try may succeed. |
| 62 | NO RUN-TIME SYSTEM NSP has not been enabled. Normal DECnet calls cannot be executed until the system manager enables NSP. |
| 66 | MISSING SPECIAL FEATURE DECnet was not installed at system generation time. The network functions cannot be executed. |

Example:

In this example, an NTCC call responds to a received Connect Initiate Message that was stored in data item RETBUF. The local link address is stored in bytes 5 and 6 of such an array and the data item LLA redefines these bytes to be referenced by that name. The call also establishes a user link address of 15, a receive maximum of 256 bytes, and indicates that the program requests message flow control. No user message data is sent so a dummy argument is given for the starting address, with zero values for the length and offset.

(Data Division)

```
01      STAT PIC 9(4) COMP.  
01      RETBUF.  
        02      FILLER PIC X(4).  
        02      LLA PIC 9(4) COMP.  
        02      FILLER PIC X(34).  
01      ULA PIC 9(4) COMP VALUE IS 15.  
01      LLMOD PIC 9(4) COMP VALUE IS 2.  
01      RMAX PIC 9(4) COMP VALUE IS 256.  
01      Z-VAL PIC 9(4) COMP VALUE IS 0.  
        .  
        .  
        .
```

(Procedure Division)

```
      CALL "NTCC" USING STAT, ULA, LLA, LLMOD,  
          RETBUF, Z-VAL, Z-VAL.
```

7.6.4 Send Connect Reject Message (NTCR)

(DECnet)

The NTCR call rejects a remote program's request for a logical link. The link is identified by the local link address established by the local NSP. Up to 16 bytes of user message data can be sent to the remote program as part of the Connect Reject Message.

COBOL Call Format:

CALL "NTCR" USING *errval*, *lla*, *reason*, *buffer*, *buflen*, *bufoff*

Argument Descriptions:

errval

The *errval* argument is a one-word COMP data item. It is set to zero if the NTCR call succeeds. If an error occurs, *errval* is set to one of the error codes listed under "Possible Errors" at the end of this section.

lla

The local link address (*lla*) is a one-word COMP data item that identifies the link being rejected. It is a number assigned by the local NSP when it received the Connect Initiate Message requesting the link. The local link address is passed to the local program in bytes 5-6 of a received Connect Initiate Message. (Section 7.7.2 describes the information returned for a received Connect Initiate Message.)

reason

DECnet/E NSP allows one of three reason code values for this one-word COMP data item: 0, 34, or 36.

Most user programs will simply use 0. If any reason is supplied for rejecting the connection, it is passed in the 16-byte user data area.

The values 34 and 36 are allowed so that DECnet/E support programs (such as FAL and NFT) can reject unauthorized or improper requests for access to local RSTS/E files in a manner agreed upon as part of the overall DECnet design. There is no reason for a user program to use these codes unless it talks to programs that interpret them. All DECnet programs that are concerned with protection against unauthorized access are designed to recognize these codes. These programs take appropriate action, such as displaying an error message to a user, when they receive such a Connect Reject Message.

The value 34 indicates that the user ID and password subfields in the Connect Data Block of the remote program's Connect Initiate Message do not correspond to any valid user known to the local system.

The value 36 indicates that, while the user ID and password subfields are acceptable, the account subfield is not. For example, the specified user may not be authorized to use that billing account or the account may be exhausted.

This argument is a data-name referring to the first character position (byte) of a buffer containing the optional user message data. The data-name can refer to a single data item or an aggregate of data items (for example, a level 01 group name). Any and all PICTURE categories can be used in the fields referenced by data-name. The message data can be offset from the beginning of the buffer (see *bufoff*).

This one-word COMP data item defines the number of bytes of user data to be sent: 0 to 16.

This one-word COMP data item defines an offset, in bytes, from the beginning of the buffer. For example, if *bufoff* = 20, the user data begins in byte 21 of the buffer.



The extended COBOL interface makes no checks for inconsistencies in the *buffer*, *buflen*, and *buffoff* values specified. A program could, for example, specify a *buflen* of 16 and a *buffoff* of -100. Sixteen bytes beginning at *buffer*-100 would be sent.

| errval | Error Text and Meaning |
|---------------|-------------------------------|
|---------------|-------------------------------|

1 **BAD DIRECTORY ON DEVICE**

Network operations have been terminated due to an internal error. Notify the system manager.

5 **CAN'T FIND FILE OR ACCOUNT**

The local link address (*lla*) does not correspond to any known logical link for the calling program. Either the *lla* is incorrect or the originating connect request has been cancelled. In the latter case, a Link Abort Message has already been queued for this link. The *lla* field in the received Link Abort Message identifies the link, since no *ula* was ever established.

10 **PROTECTION VIOLATION**

Some procedural error has occurred. Sending a Connect Reject Message for a link that is not waiting for confirmation will cause this error. For example, sending a Connect Reject Message for a link that is already established will return this error.

Extended COBOL Interface Call Formats 7-37

errval Error Text and Meaning

18 ILLEGAL SYS() USAGE

One of two possibilities:

1. The calling program is not a declared receiver. A Declare Receiver call must be issued before this call is given.
2. The reason code supplied with the Connect Reject Message was not 0, 34, or 36.

31 ILLEGAL BYTE COUNT FOR I/O

The length field passed in *buflen* is invalid. The length of the buffer must be between 0 and 16 bytes for a Connect Reject Message.

32 NO BUFFER SPACE AVAILABLE

System buffers are not currently available to store this message. A later try may succeed.

62 NO RUN-TIME SYSTEM

NSP has not been enabled. Normal DECnet calls cannot be executed until the system manager enables NSP.

66 MISSING SPECIAL FEATURE

DECnet was not installed at system generation time. The network functions cannot be executed.

Example:

In the following example, an NTCR call responds to a received Connect Initiate Message stored in the receiving buffer RETBUF. A reason code of 0 is supplied. The local program notifies the remote program of the reason for rejection with user data.

(Data Division)

```
01        STAT PIC 9(4) COMP.
01        RETBUF.
          03        FILLER PIC X(4).
          03        LLA PIC 9(4) COMP.
          03        FILLER PIC X(34).
01        BUFFER PIC X(8) VALUE IS "TOO LATE".
01        BUFLN PIC 9(4) COMP VALUE IS 8.
01        BUFOFF PIC 9(4) COMP VALUE IS 0.
01        REASON PIC 9(4) COMP VALUE IS 0.
          .
          .
          .
```

(Procedure Division)

```
CALL "NTCR" USING STAT, LLA, REASON, BUFFER,
          BUFLN, BUFOFF.
```


7.6.5 Send Network Data Message (NTDM)

(DECnet)

The NTDM call transmits user data — a *segment* as defined in Chapter 6 — over an established logical link. A flag byte indicates whether or not the segment is the end of a logical message. If the remote program requested segment or message flow control when the link was established, requests for data must be outstanding or this call will result in an error.

COBOL Call Format:

CALL "NTDM" USING *errval*, *retbuf*, *ula*, *dmflgs*, *buffer*,
buflen, *bufoff*

Argument Descriptions:

errval

The *errval* argument is a one-word COMP data item. It is set to zero if the NTDM call succeeds. If an error occurs, *errval* is set to one of the error codes listed under "Possible Errors" at the end of this section.

retbuf

This argument is a data-name referring to the first character position (byte) of a buffer to receive link status information returned by the call. This argument must be specified regardless of whether or not link status information is requested and should refer to a buffer at least 40 bytes long.

ula

The *ula* is a one-word COMP data item giving the user link address (1 to 255) chosen by the calling program in the Connect Initiate (NTCI) or Connect Confirm (NTCC) Message sent during the connection sequence.

dmflgs

The *dmflgs* argument is a one-word COMP data item. It can be coded as the sum of three bit values:

$$dmflgs = rls + eom + bom$$

The low-order two bits of the *dmflgs* argument (*eom* + *bom*) indicate whether the segment being sent in this Network Data Message is the first, last, middle, or sole segment of a logical message. If message flow control was requested by the remote program when the link was established, NSP uses these two bits to regulate the transmission of logical messages as the remote program requests them. (See Chapter 6 for a full discussion of flow control.)

| <i>eom</i> | <i>bom</i> | |
|------------|------------|-------------------------------------|
| 0 | 0 | Middle segment of a logical message |
| 0 | 1 | First segment of a logical message |
| 2 | 0 | Last segment of a logical message |
| 2 | 1 | Sole segment of a logical message |

DECnet/E NSP uses the *eom* bit to count logical messages transmitted to the remote program. The *bom* bit is not used in maintaining the logical message request count but is simply passed to the remote system and program.

NSP also uses the *eom* flag to determine whether it will permit a disconnect of the logical link. If a disconnect is requested but the last Network Data Message sent did not have the *eom* flag set, NSP assumes that the last message has not been completely sent and will not permit the disconnect. This is always done, irrespective of the flow control option selected by the receiver.

Some DECnet implementations require that the *eom* flag be set to terminate an asynchronous I/O operation. It is suggested, therefore, that the *eom* flag be set whenever segment flow control or no flow control is requested by the remote program.

The *rls* bit is used to request that link status information be returned to the *retbuf* buffer when the call has completed successfully.

rls = 128 Link status information is returned.
rls = 0 No link status information is returned.

If link status information is requested, it is returned in the same format as a received Link Service Message (see Section 7.7.7). If no link status information is requested, random data is returned.

buffer


This argument is a data-name referring to the first character position (byte) of a buffer containing the user message data. The data-name can refer to a single data item or an aggregate of data items (for example, a level 01 group name). Any and all PICTURE categories can be used in the fields referenced by data-name. The data can be offset from the beginning of the buffer (see *bufoff*).

buflen

This one-word COMP data item defines the number of bytes to be sent. The maximum amount of data that can be transferred over the logical link is established by the remote program or by the remote NSP. It is passed on to the calling program as a transmit maximum (the *tmax* value in a received Connect Initiate or Connect Confirm Message, Sections 7.7.2 and 7.7.3). Thus, the value of *buflen* can range from 1 through *tmax*.

bufoff

This one-word COMP data item defines an offset, in bytes, from the start of the buffer. For example, if *bufoff* = 20, the user data begins in byte 21 of the buffer.



NOTE

The extended COBOL interface makes no checks for inconsistencies in the *buffer*, *buflen*, and *bufoff* values specified. A program could, for example, specify a *buflen* of 100 and a *bufoff* of 20. One hundred bytes beginning at *buffer*+20 would be sent.

Possible Errors:

| <i>errval</i> | Error Text and Meaning |
|---------------|--|
| 1 | BAD DIRECTORY ON DEVICE Network operations have been terminated due to an internal error. Notify the system manager. |
| 4 | NO ROOM FOR USER ON DEVICE The Data Message transmit queue for this logical link is full. NSP is waiting for acknowledgment from the remote system for previously sent Network Data Messages. No more can be sent until at least one of the outstanding messages in the transmit queue is acknowledged. This condition is temporary and no Link Service Message notification occurs when the condition clears. The program should simply retry the send after a short delay of one to five seconds. (See Chapter 6 for a full discussion of flow control.) |
| 5 | CAN'T FIND FILE OR ACCOUNT The user link address (<i>ula</i>) specified in the call does not correspond to any known logical link for the calling program. Either the <i>ula</i> is incorrect or the logical link has been disconnected. In the latter case, a Disconnect or Link Abort Message has already been queued for this link. |
| 10 | PROTECTION VIOLATION Some procedural error has occurred. Sending a Network Data Message over a logical link that has not yet been confirmed will cause this error. |
| 18 | ILLEGAL SYS() USAGE One of two possibilities: <ol style="list-style-type: none">1. The calling program is not a declared receiver. A Declare Receiver call must be issued before this call is given.2. The <i>ula</i> specified is zero. It must be within the range 1 to 255. |
| 19 | DISK BLOCK IS INTERLOCKED Either the remote system has inhibited transmission on this link because of a backpressure condition or there are no outstanding requests for segments or logical messages from the remote program (assuming it requested segment or message flow control for the link). The calling program will be notified with a Link Service Message when the condition clears, as described in Chapter 6. |

(continued on next page)

errval Error Text and Meaning

- | | |
|----|---|
| 31 | ILLEGAL BYTE COUNT FOR I/O The length field passed in <i>buflen</i> is invalid. The length of the buffer must be between 1 and the transmit maximum (<i>tmax</i>) established by the remote program or system in its Connect Initiate or Connect Confirm Message for this logical link. |
| 32 | NO BUFFER SPACE AVAILABLE System buffers are not currently available to store this message. A later try may succeed. |
| 62 | NO RUN-TIME SYSTEM NSP has not been enabled. Normal DECnet calls cannot be executed until the system manager enables NSP. |
| 66 | MISSING SPECIAL FEATURE DECnet was not installed at system generation time. The network functions cannot be executed. |

Example:

The following example sends a Network Data Message over a logical link. The user link address is 20. The value for DMFLGS identifies the message data as the sole segment of a logical message and requests the return of link status information. The buffer for returned data is set up as a 40-byte data item named RETBUF.

The call sends 400 bytes of data from BUFFER. After the call, the program checks for an error. If none occurred, the program then determines if there are any pending messages. If so, the program branches to receive and process messages. If not, the program simply continues.

(Data Division)

```

01      RETBUF.
03      FILLER PIC X(2).
03      CODE-ULA PIC 9(4) COMP.
03      FILLER PIC X(4).
03      REMLen PIC 9(4) COMP.
03      FILLER PIC X(2).
03      MESLEN PIC 9(4) COMP.
03      FILLER PIC X(6).
03      L-AND-R-STAT PIC 9(4) COMP.
03      L-AND-R-DRQ PIC 9(4) COMP.
03      L-AND-R-IRQ PIC 9(4) COMP.
03      DTM-AND-DTC PIC 9(4) COMP.
03      LTM-AND-LTC PIC 9(4) COMP.
03      MMAX-AND-MCNT PIC 9(4) COMP.
03      MCNT REDEFINES MMAX-AND-MCNT.
05      MMAX PIC X.
05      MCNTX PIC X.
03      MULA PIC 9(4) COMP.
03      FILLER PIC X(16).
01      MCNT-GOOD PIC 9(4) COMP.
01      WORK-WORD REDEFINES MCNT-GOOD.
03      LOW PIC X.
03      HIGH PIC X.
01      STAT PIC 9(4) COMP.
01      BUFFER PIC X(400).
01      BUFLen PIC 9(4) COMP VALUE IS 400.
01      BUFOFF PIC 9(4) COMP VALUE IS 0.
01      ULA PIC 9(4) COMP VALUE IS 20.
01      EOM PIC 9(4) COMP VALUE IS 1.
01      BOM PIC 9(4) COMP VALUE IS 2.
01      RLS PIC 9(4) COMP VALUE IS 128.
01      DMFLGS PIC 9(4).

```

.
.
.

(Procedure Division)

```

ADD EOM, BOM, RLS GIVING DMFLGS.
CALL "NTDM" USING STAT, RETBUF, ULA, DMFLGS,
      BUFFER, BUFLen, BUFOFF.
IF STAT NOT EQUAL 0 GO TO ERR-PROC.
MOVE 0 TO MCNT-GOOD.
MOVE MCNTX TO LOW.
IF MCNT-GOOD NOT EQUAL 0 GO TO RECV-MSG.

```

7.6.6 Send Interrupt Message (NTIN)

(DECnet)

The NTIN call transmits an Interrupt Message to a remote program over an established logical link. All DECnet systems are designed to deliver Interrupt Messages ahead of other messages. If the remote system is a DECnet/E system, the Interrupt Message will be placed at the head of the pending message queue, behind the first message (since it can already be partly received) and behind any other pending Interrupt Messages queued for the program. Interrupt messages are subject to flow control, as described in Chapter 6.

COBOL Call Format:

CALL "NTIN" USING *errval*, *retbuf*, *ula*, *inflgs*, *buffer*,
buflen, *bufoff*

Argument Descriptions:

errval

The *errval* argument is a one-word COMP data item. It is set to zero if the NTIN call succeeds. If an error occurs, *errval* is set to one of the error codes listed under "Possible Errors" at the end of this section.

retbuf

This argument is a data-name referring to the first character position (byte) of a buffer to receive link status information returned by the call. This argument must be specified regardless of whether or not link status information is requested and should refer to a buffer at least 40 bytes long.

ula

The *ula* is a one-word COMP data item specifying the user link address (1 to 255) chosen by the calling program in the Connect Initiate (NTCI) or Connect Confirm (NTCC) Message sent during the connection sequence.

inflgs

The *inflgs* argument indicates whether link status information is to be returned to *retbuf* when the call has completed successfully. This field should be coded and used as follows:

inflgs = *rls*

where

rls = 128 Link status information is returned.

rls = 0 No link status information is returned.

If link status information is requested, it is returned in the same format as a received Link Service Message (see Section 7.7.7). If no link status information is requested, random data is returned.


This argument is a data-name referring to the first character position (byte) of a buffer containing optional user message data. The data-name can refer to a single data item or an aggregate of data items (for example, a level 01 group name). Any and all PICTURE categories can be used in the fields referenced by data-name. The data can be offset from the beginning of the buffer (see *bufoff*).

buflen

This one-word COMP data item defines the number of bytes of user data to be sent: 0 to 16.

butoff

This one-word COMP data item defines an offset, in bytes, from the beginning of the buffer. For example, if *bufoff* = 20, the user data begins in byte 21 of the buffer.



NOTE

The extended COBOL interface makes no checks for inconsistencies in the *buffer*, *buflen*, and *buffoff* values specified. A program could, for example, specify a *buflen* of 16 and a *buffoff* of -100. Sixteen bytes beginning at *buffer*-100 would be sent.

Possible Errors:

| error | Error Text and Meaning |
|--------------|--|
| 1 | <p>BAD DIRECTORY ON DEVICE</p> <p>Network operations have been terminated due to an internal error. Notify the system manager.</p> |
| 4 | <p>NO ROOM FOR USER ON DEVICE</p> <p>The Interrupt/Link Service transmit queue for this logical link is full. NSP is waiting for acknowledgment from the remote system for previously sent Interrupt or Link Service Messages. No Interrupt or Link Service Messages can be sent until at least one of the outstanding messages in the transmit queue is acknowledged. This condition is temporary and no Link Service Message notification occurs when the condition clears. The program should simply retry the send after a short delay of one to five seconds. (See Chapter 6 for a full discussion of flow control.)</p> |

(continued on next page)

| | |
|----|---|
| 5 | CAN'T FIND FILE OR ACCOUNT The local link address (<i>lla</i>) does not correspond to any known logical link for the calling program. Either the <i>lla</i> is incorrect or the originating connect request has been cancelled. In the latter case, a Disconnect or Link Abort Message has already been queued for this link. |
| 10 | PROTECTION VIOLATION Some procedural error has occurred. Sending an Interrupt Message over a logical link that has not yet been confirmed will cause this error. |
| 18 | ILLEGAL SYS() USAGE One of two possibilities: <ol style="list-style-type: none">1. The calling program is not a declared receiver. A Declare Receiver call must be issued before this call is given.2. The user link address (<i>ula</i>) given was zero. Valid values range from 1 to 255. |
| 19 | DISK BLOCK IS INTERLOCKED The remote NSP has prohibited transmission of Interrupt Messages over this logical link according to its rules for Interrupt Message flow control. A Link Service Message will be delivered to the calling program when the condition clears. (See Chapter 6 for a full discussion of flow control.) |
| 31 | ILLEGAL BYTE COUNT FOR I/O The length field passed in <i>buflen</i> is invalid. The length of the buffer must be between 0 and 16 bytes for an Interrupt. |
| 32 | NO BUFFER SPACE AVAILABLE System buffers are not currently available to store this message. A later try may succeed. |
| 62 | NO RUN-TIME SYSTEM NSP has not been enabled. Normal DECnet calls cannot be executed until the system manager enables NSP. |
| 66 | MISSING SPECIAL FEATURE DECnet was not installed at system generation time. The network functions cannot be executed. |

Example:

The following example shows a NTIN call for logical link 15. The call sends the message data "TIMEOUT" and requests no link status information. A buffer must be defined and referenced, however, so an area named DUMMY is defined and used in the call.

(Data Division)

```
01      ULA PIC 9(4) COMP VALUE IS 15.  
01      DUMMY PIC X(40).  
01      BUFFER PIC X(7) VALUE IS "TIMEOUT".  
01      BUFLen PIC 9(4) COMP VALUE IS 7.  
01      BUFOFF PIC 9(4) COMP VALUE IS 0.  
01      STAT PIC 9(4) COMP VALUE IS 0.  
01      INFLGS PIC 9(4) COMP VALUE IS 0.
```

·
·
·

(Procedure Division)

```
CALL "NTIN" USING STAT, DUMMY, ULA, INFLGS,  
                BUFFER, BUFLen, BUFOFF.
```

7.6.7 Send Link Service Message (NTLS)

(DECnet)

The NTLS call can be used to: (1) request data from a remote program over a flow-controlled logical link (Chapter 6), (2) reenale incoming Interrupt Messages over a logical link (Chapter 6), or (3) obtain status information for the link. No user data is allowed with a Link Service Message.

COBOL Call Format:

CALL "NTLS" USING *errval*, *retbuf*, *ula*, *lsflgs*, *drcnt*, *ircnt*

Argument Descriptions:

errval

The *errval* argument is a one-word COMP data item. It is set to zero if the NTLS call succeeds. If an error occurs, *errval* is set to one of the error codes listed under "Possible Errors" at the end of this section.

retbuf

This argument is a data-name referring to the first character position (byte) of a buffer to receive link status information returned by the call. The argument must be specified in the call regardless of whether or not link status information is requested and should refer to a buffer at least 40 bytes long.

ula

The *ula* is a one-word COMP data item giving the user link address (1 to 255) chosen by the calling program in the Connect Initiate (NTCI) or Connect Confirm (NTCC) Message sent during the connection sequence.

lsflgs

The link service flags argument is a one-word COMP data item coded as the sum of two values: $lsflgs = rls + opr$. The low-order two bits (*opr*) indicate the purpose of the call.

- opr* = 0 Indicates the call is requesting segments or logical messages from the remote program. In this case, the *drcnt* argument is interpreted as the number of segments or logical messages being requested. This form of the call can be used only when the calling program requested message or segment flow control in its Connect Initiate or Connect Confirm Message for this logical link. If no flow control was requested, an NTLS call with *opr* = 0 will fail with *errval* = 18.
- = 1 Indicates the call is to reenale incoming Interrupt Messages for this link. In this case, the *ircnt* argument is interpreted as the number of Interrupt Messages requested (always 1).
- = 2 Indicates the call is to obtain status information only. In this case, the *drcnt* and *ircnt* arguments are ignored.

The *rls* bit can be used when *opr* = 0 or 1 to request that link status information be returned to the *retbuf* buffer when the call completes successfully.

rls = 128 Link status information is returned.
= 0 No link status information is returned.

Link status information is always returned when *opr* = 2. If link status information is requested, it is returned in the same format as a received Link Service Message (see Section 7.7.7). If no link status information is requested, random data is returned.

drcnt

The data request count (*drcnt*) is meaningful only when *opr* = 0 and indicates the maximum number of segments or logical messages being requested from the remote program. Depending on the flow control option in effect for this end of the logical link, the value is interpreted as either a segment counter (segment flow control) or a logical message counter (message flow control). In either case, it is an incremental count, added to the counter of outstanding requests for segments or logical messages kept by the remote and local NSP (see Chapter 6). An error is returned if the *drcnt* specified increases the counter to more than 127.

ircnt

The interrupt request count (*ircnt*) is a one-word COMP data item. It is meaningful only when *opr* = 1 and indicates the number of Interrupt Messages being requested. DECnet/E allows only one interrupt request outstanding for a logical link at any given time. Hence, *ircnt* must always be 1 when *opr* = 1 and has the effect of reenabling interrupts on the logical link. Remember that Interrupt Messages can be reenabled only when the previous Interrupt Message for this logical link, if any, has been received from the pending message queue. If it has not, the call will fail with an error (*errval* = 19).

Possible Errors:

| <i>errval</i> | Error Text and Meaning |
|----------------------|--|
| 1 | BAD DIRECTORY ON DEVICE Network operations have been terminated due to an internal error. Notify the system manager. |
| 4 | NO ROOM FOR USER ON DEVICE The Interrupt/Link Service transmit queue for this logical link is full. NSP is waiting for acknowledgment from the remote system for previously sent Interrupt or Link Service Messages. No Interrupt or Link Service Messages can be sent until at least one of the outstanding messages in the transmit queue is acknowledged. This condition is temporary and will not cause a Link Service Message to be queued when the condition clears. The program should simply retry the send after a short delay of one to five seconds. (See Chapter 6 for a full discussion of flow control.) |

(continued on next page)

errval Error Text and Meaning

| | |
|----|---|
| 5 | CAN'T FIND FILE OR ACCOUNT The user link address (<i>ula</i>) specified in the call does not correspond to any known logical link for the calling program. Either the <i>ula</i> is incorrect or the logical link has been disconnected. In the latter case, a Disconnect or Link Abort Message has already been queued for this link. |
| 10 | PROTECTION VIOLATION Some procedural error has occurred. Sending a Link Service Message over a logical link that has not yet been confirmed will cause this error. |
| 18 | ILLEGAL SYS() USAGE One of several possibilities: <ol style="list-style-type: none">1. The calling program is not a declared receiver. A Declare Receiver call must be issued before this call is given.2. The <i>opr</i> value is not 0, 1, or 2.3. The <i>opr</i> value is 0, but the calling program did not request flow control for the logical link in its Connect Initiate (NTCI) or Connect Confirm (NTCC) Message.4. The data request count (<i>drcnt</i>) increased the count of outstanding requests to more than 127.5. The <i>opr</i> value is 1, indicating a request for an Interrupt Message, but the <i>ircnt</i> argument does not equal 1.6. The <i>opr</i> value is 1, indicating a request for an Interrupt Message, but the outstanding requests counter for Interrupt Messages already equals 1.7. The <i>ula</i> specified is zero. It must be within the range 1 to 255. |
| 19 | DISK BLOCK IS INTERLOCKED The <i>opr</i> value is 1, indicating a request for an Interrupt Message, but there is already one queued for this logical link. See Chapter 6 for details on flow control. |
| 32 | NO BUFFER SPACE AVAILABLE System buffers are not currently available to store this message. A later try may succeed. |
| 62 | NO RUN-TIME SYSTEM NSP has not been enabled. Normal DECnet calls cannot be executed until the system manager enables NSP. |
| 66 | MISSING SPECIAL FEATURE DECnet was not installed at system generation time. The network functions cannot be executed. |

Example:

The following example of the NTLS call requests five segments from the remote program over logical link 23. (Assume that the calling program requested segment flow control in its NTCI or NTCC call for this link).

(Data Division)

```
01      STAT PIC 9(4) COMP.  
01      DUMMY PIC X(40).  
01      OPR PIC 9(4) COMP VALUE IS 0.  
01      RLS PIC 9(4) COMP VALUE IS 128.  
01      LSFLGS PIC 9(4) COMP.  
01      ULA PIC 9(4) COMP VALUE IS 23.  
01      DRCNT PIC 9(4) COMP VALUE IS 5.  
01      IRCNT PIC 9(4) COMP VALUE IS 0.
```

.
.
.

(Procedure Division)

```
CALL "NTLS" USING STAT, DUMMY, ULA,  
                LSFLGS, DRCNT, IRCNT.
```

7.6.8 Send Disconnect Message (NTDI)

(DECnet)

The NTDI call terminates an established logical link. Messages in the pending message queue will remain but no more messages can be sent over the link. The user link address is freed and can be used for another logical link. If the last Network Data Message sent did not have the end-of-message flag set, or if any Network Data, Interrupt, or Link Service Messages are still waiting for acknowledgment in the transmit queues for this logical link, this call will fail with an error.

Successful completion of an NTDI call assures the sender that the remote system has received and acknowledged all Network Data, Interrupt, and Link Service Messages previously sent over the link. It does not, however, guarantee that the receiving program has processed the messages. The Disconnect Message is useful in "master-slave" communication where the master program only transmits data and the slave program only receives data (see Section 5.13). Otherwise, the NTDI call provides no particular advantage over an NTLA call (see Section 7.6.9) in terminating a logical link.

Up to 16 bytes of user data can be sent with the Disconnect Message.

COBOL Call Format:

CALL "NTDI" USING *errval*, *ula*, *buffer*, *buflen*, *bufoff*

Argument Descriptions:

errval

The *errval* argument is a one-word COMP data item. It is set to zero if the NTDI call succeeds. If an error occurs, *errval* is set to one of the error codes listed under "Possible Errors" at the end of this section.

ula

The *ula* is a one-word COMP data item giving the user link address (1 to 255) chosen by the calling program in the Connect Initiate (NTCI) or Connect Confirm (NTCC) Message sent during the connection sequence.


buffer

This argument is a data-name referring to the first character position (byte) of a buffer containing the optional user message data. The data-name can refer to a single data item or an aggregate of data items (for example, a level 01 group name). Any and all PICTURE categories can be used in the fields referenced by data-name. The data can be offset from the beginning of the buffer (see *bufoff*).

buflen

This one-word COMP data item defines the number of bytes of user data to be sent: 0 to 16.

This one-word COMP data item defines an offset, in bytes, from the beginning of the buffer. For example, if *bufoff* = 20, the user data begins in byte 21 of the buffer.



NOTE

The extended COBOL interface makes no checks for inconsistencies in the *buffer*, *buflen*, and *bufoff* values specified. If, for example, a program specifies a buffer 20 bytes long with an offset of 10 and a *buflen* of 16, 16 bytes will be sent: the last 10 bytes of the buffer, and 6 bytes of whatever is beyond the buffer in memory.

Possible Errors:

| error | Error Text and Meaning |
|-------|---|
| 1 | <p>BAD DIRECTORY ON DEVICE</p> <p>Network operations have been terminated due to an internal error. Notify the system manager.</p> |
| 4 | <p>NO ROOM FOR USER ON DEVICE</p> <p>There are outstanding unacknowledged messages on either the Data Message transmit queue or the Interrupt/Link Service transmit queue. All messages previously sent over the logical link must be acknowledged before the Disconnect Message can be sent. (If an immediate unconditional termination of the logical link is desired, use the NTLA call.)</p> |
| 5 | <p>CAN'T FIND FILE OR ACCOUNT</p> <p>The user link address (<i>ula</i>) specified in the call does not correspond to any known logical link for the calling program. Either the <i>ula</i> is incorrect or the logical link has been disconnected. In the latter case, a Disconnect or Link Abort Message has already been queued for this link.</p> |
| 10 | <p>PROTECTION VIOLATION</p> <p>Some procedural error has occurred. Sending a Disconnect Message over a logical link that has not yet been confirmed will cause this error.</p> |
| 18 | <p>ILLEGAL SYS() USAGE</p> <p>One of two possibilities:</p> <ol style="list-style-type: none"> 1. The calling program is not a declared receiver. A Declare Receiver call must be issued before this call is given. 2. The <i>ula</i> specified is zero. It must be within the range 1 to 255. |

(continued on next page)

errval Error Text and Meaning

| | |
|----|---|
| 19 | DISK BLOCK IS INTERLOCKED The last segment sent over this logical link did not have the <i>eom</i> flag set. NSP will not permit the link to be disconnected unless all messages have been completely transmitted. (Note that this test of the <i>eom</i> flag is done irregardless of the flow control option chosen by the remote program.) |
| 31 | ILLEGAL BYTE COUNT FOR I/O The length field passed in <i>buflen</i> is invalid. The length of the buffer must be between 0 and 16 for a Disconnect Message. |
| 32 | NO BUFFER SPACE AVAILABLE System buffers are not currently available to store this message. A later try may succeed. |
| 62 | NO RUN-TIME SYSTEM NSP has not been enabled. Normal DECnet calls cannot be executed until the system manager enables NSP. |
| 66 | MISSING SPECIAL FEATURE DECnet was not installed at system generation time. The network functions cannot be executed. |

Example:

The following example disconnects logical link 255 and includes user data explaining the reason for the disconnect. (The remote program would have to be coded to recognize this message and take appropriate action.)

(Data Division)

```
01      ULA PIC 9(4) COMP VALUE IS 255.  
01      STAT PIC 9(4) COMP.  
01      BUFFER PIC X(6) VALUE IS "NOMORE".  
01      BUFLN PIC 9(4) COMP VALUE IS 6.  
01      BUFOFF PIC 9(4) COMP VALUE IS 0.
```

;
;
;

(Procedure Division)

```
CALL "NTDI" USING STAT, ULA, BUFFER, BUFLN, BUFOFF.
```


7.6.9 Send Link Abort Message (NTLA)

(DECnet)

The NTLA call terminates a logical link immediately. If there are still messages in the transmit queues waiting for acknowledgment from the remote NSP, they are discarded. They may or may not have been received by the remote system (see Section 5.14).

Up to 16 bytes of user data can be sent with the message. If the NTLA call is issued for an established logical link, the remote system is notified that the link has been broken and the user data is delivered to the remote program. If the NTLA call is issued for a logical link awaiting confirmation from the remote program, the user data does not reach the remote program (see Section 5.14).

COBOL Call Format:

CALL "NTLA" USING *errval*, *ula*, *buffer*, *buflen*, *bufoff*

Argument Descriptions:

errval

The *errval* argument is a one-word COMP data item. It is set to zero if the NTLA call succeeds. If an error occurs, *errval* is set to one of the error codes listed under "Possible Errors" at the end of this section.

ula

The *ula* is a one-word COMP data item giving the user link address (1 to 255) chosen by the calling program in the Connect Initiate (NTCI) or Connect Confirm (NTCC) Message sent during the connection sequence.

buffer

This argument is a data-name referring to the first character position (byte) of a buffer containing the optional user message data. The data-name can refer to a single data item or an aggregate of data items (for example, a level 01 group name). Any and all PICTURE categories can be used in the fields referenced by data-name. The data can be offset from the beginning of the buffer (see *bufoff*).

buflen

This one-word COMP data item defines the number of bytes of user data to be sent: 0 to 16.

bufoff

This one-word COMP data item defines an offset, in bytes, from the beginning of the buffer. For example, if *bufoff* = 20, the user data begins in byte 21 of the buffer.

Example:

The following example aborts logical link 17. No user data is sent.

(Data Division)

```
01      STAT PIC 9(4) COMP.  
01      ULA PIC 9(4) COMP VALUE IS 17.  
01      DUMMY PIC X.  
01      BUFLN PIC 9(4) COMP VALUE IS 0.  
      .  
      .  
      .
```

(Procedure Division)

```
CALL "NTLA" USING STAT, ULA, DUMMY, BUFLN, BUFLN.
```

7.7 Receive (MRCV)

(Local and DECnet)

The MRCV call retrieves a message from the calling program's pending message queue. It returns control parameters and user data, if any, to buffers defined in the call. The Connect Data Block for Connect Initiate Messages is also delivered to the buffer that receives the user data. The message data can be retrieved all at once or in portions, and portions can be discarded. The formats in Sections 7.7.1 – 7.7.9 show what control information is returned for the various types of messages.

A Receive call can be selected to get (1) the first message in the queue, (2) the first message in the queue from either a network or a local sender, or (3) the first message in the queue for a particular logical link. A sleep flag and associated timer can be set to indicate that the calling program is to wait if there are no appropriate messages pending in the queue.

COBOL Call Format:

CALL "MRCV" USING *errval*, *retbuf*, *rmod*, *sndr*, *qual*, *buffer*,
buflen, *bufoff*, *slptime*

Argument Descriptions:

errval

The *errval* argument is a one-word COMP data item. It is set to zero if the MRCV call succeeds. If an error occurs, *errval* is set to one of the error codes listed under "Possible Errors" at the end of this section.

retbuf

This argument is a data-name referring to the first character position (byte) of a buffer to contain the control information identifying the type and characteristics of the message received. This argument must refer to a buffer at least 40 bytes long. The format of the data returned depends on the message type, as described in Sections 7.7.1 – 7.7.9.

rmod

The receive modifier (*rmod*) argument is a one-word COMP data item consisting of four bit values: $s + t + l + n$.

- s The sleep bit (s) defines what is to be done if there is no appropriate message in the queue for the Receive call to deliver — that is, a message of the type requested with the l , n , *sndr*, and *qual* arguments. If $s = 1$, the program will sleep — suspend execution — until something happens. The *slptime* argument regulates the time that the program will sleep.

If the s bit is cleared to 0, the program will not sleep. If there are no appropriate messages in the queue the call terminates immediately with *errval* = 5.

t The truncate bit (*t*) is set to 2 to discard any excess information that will not fit in the buffer specified in this call. Information could be left over if the value of *buflen* is smaller than the amount of user data in the message.

A 0 value for *t* indicates that such excess information is to be kept for retrieval on later Receive calls.

l The local sender bit (*l*) is set to 4 to select a message from a local sender. The selection can be further qualified with the *sndr* and *qual* arguments.

If *l* = 0, local selection is not requested. If both *l* and *n* are zero, the first message in the queue is returned.

n The network sender bit (*n*) is set to 8 to select a message from a network sender. The selection can be further qualified with the *sndr* and *qual* arguments.

If *n* = 0, network selection is not requested. If both *l* and *n* = 0, the first message pending in the queue is returned. If *l* = 4, and *n* = 8, local selection prevails.

sndr

The sender selection argument (*sndr*) is a one-word COMP data item. It selects a particular local sender or a particular logical link for this Receive call.

If *l* = 4 (local selection) and *sndr* is nonzero, *sndr* is interpreted as a job number times two. The first message on the queue from that particular local job is retrieved.

If *l* = 0, and *n* = 8 (network selection), a nonzero value for *sndr* is interpreted as a user link address. The first message on the queue from that logical link is delivered.

If *sndr* = 0 when either *l* = 4 or *n* = 8, the *qual* parameter has special meaning. The *sndr* and *qual* arguments are ignored if both *l* and *n* are zero.

qual

The sender selection qualifier (*qual*) argument is a one-word COMP data item. It is normally zero for user applications.

If *l* = 4 (local selection) and *sndr* = 0, any nonzero value for *qual* is a special-case Receive call, requesting a message from the system (job number 0). This special case is used by the RSTS/E utility ERRCPY, which processes messages from the monitor error logging routines.

The *qual* argument is ignored if both *l* and *n* = 0, or if *sndr* ≠ 0.

Table 7-4 summarizes the relationships between the selection bits of *rmod* and the *sndr* and *qual* arguments.

Table 7-4: Sender Selection Summary

| <i>rmod</i> | | | | |
|-------------|----------|-------------|-------------|--|
| <i>n</i> | <i>l</i> | <i>sndr</i> | <i>qual</i> | Result |
| 0 | 0 | n/a | n/a | The <i>sndr</i> and <i>qual</i> values are ignored. The MRCV call returns the first message in the pending message queue. |
| n/a | 4 | 0 | 0 | Selects the first Local Data Message in the queue. |
| | | 0 | nonzero | Selects job 0. This combination is used by the error logging programs to select messages from monitor error logging routines. |
| | | nonzero | n/a | Selects Local Data Messages by job number. The <i>sndr</i> argument is interpreted as a job number times 2. Only messages from that job are delivered on this MRCV call. |
| 8 | 0 | 0 | n/a | Selects the first network message (anything other than a Local Data Message). |
| | | nonzero | n/a | Selects network messages from a particular logical link. The <i>sndr</i> argument is interpreted as a user link address. Only network messages from the designated logical link are delivered on this MRCV call. |

buffer

This argument is a data-name referring to the first character position (byte) of the buffer to which user message data is to be delivered. The data-name can refer to a single data item or an aggregate of data items (for example, a level 01 group name). Any and all PICTURE categories can be used in fields referenced by data-name. The format should, of course, reflect the type of data sent from the remote program.

The data can be offset from the beginning of the buffer (see *bufoff*).

buflen

This one-word COMP data item defines the maximum number of bytes to be delivered to the buffer on this Receive call.


If a program has limited buffer space, the *buflen* value can be set to process a large segment in small pieces. If truncation is not requested (*t* bit of *rmod* = 0) and the amount of data in the message is greater than the amount indicated by *buflen*, then only the amount indicated by *buflen* will be returned and the rest will be saved for retrieval with later Receive calls.

The usual sequence in this case is to issue Receive calls until the number of bytes remaining in the pending message goes to zero. This point can be deter-

mined by examining bytes 9-10 in the data returned to the *retbuf* buffer by the Receive call. The number of bytes actually returned to the buffer is returned in bytes 13-14 in the *retbuf* buffer. (See the formats in Sections 7.7.1 - 7.7.9.)

bufoff

This one-word COMP data item defines an offset, in bytes, from the beginning of the buffer. For example, if *bufoff* = 20, the user message data is delivered beginning in byte 21 of the buffer.



NOTE

The extended COBOL interface makes no checks for inconsistencies in the *buffer*, *buflen*, and *bufoff* values specified. A program could, for example, specify a *buflen* of 100 and a *bufoff* of 20. Up to 100 bytes beginning at *buffer*+20 could be delivered on the Receive call.

slptime

The *slptime* argument is a one-word COMP data item. It is significant only when the *s* bit of *rmod* = 1. In that case, a positive value for *slptime* defines the length of time, in seconds, that a program is to sleep when no appropriate message is in the queue. The program will sleep until:

1. The sleep timer expires.
2. Any new message is placed in the queue (not just an appropriate one).
3. A delimiter is typed on a terminal opened by, or assigned to, the job.
4. Log-ins are disabled. (This could occur if the system is being shut down).
5. A state change occurs on a pseudo-keyboard assigned to the job. (The job has printed output for the controlling job to read or has entered an input wait state.)
6. The job itself has opened the XM: (DMC11 driver) and a message is received for the job.

In all cases, the program is awakened with an error (*errval* = 5) but is not passed a message. To obtain a pending message, the program must execute another Receive call. Since the program can be awakened by any of the conditions listed previously, a check for pending messages can be made by executing a MRCV call without a sleep.

Possible Errors:

| errval | Error Text and Meaning |
|---------------|---|
| 5 | CAN'T FIND FILE OR ACCOUNT For a Receive call without sleep (<i>s</i> bit in <i>rmod</i> = 0), this error indicates that no appropriate messages are pending. For a Receive call with sleep (<i>s</i> bit in <i>rmod</i> = 1), this error is returned when the program is awakened from the sleep. The program must execute another MRCV call to retrieve any pending messages. |
| 18 | ILLEGAL SYS() USAGE No Receiver ID block exists for the calling program. Before any Receive call can succeed, a Declare Receiver (MDCL) call must be executed. |

Example:

In the following example, the first Receive call is issued for network messages with the sleep bit set and the timer set to five seconds. The truncate bit is 0 and the length argument is 0, so all the information that is available is saved for the second Receive call. The test for `ERRVAL = 5` after the first Receive call, with sleep, ensures that a message is waiting when the second Receive call is issued. The "BYTSGN" subprogram is the sample program shown in Section 7.1 to convert signed one-byte binary data items to signed one-word COMP data items. The GO TO then checks the value identifying the type of message and branches accordingly.

(Data Division)

```

01      S PIC 9(4) COMP VALUE IS 1.
01      N PIC 9(4) COMP VALUE IS 8.
01      RMOD PIC 9(4) COMP.
01      STAT PIC 9(4) COMP.
01      SLPTIME PIC 9(4) COMP VALUE IS 5.
01      Z-VAL PIC 9(4) COMP VALUE IS 0.
01      RETBUF.
03      RETBUF-CHAR PIC X OCCURS 40 TIMES.
01      BUFFER PIC X(256).
01      BUFLen PIC 9(4) COMP VALUE IS 256.
01      ABC PIC X.
01      DEF PIC 9(4) COMP.
      .
      .
      .

```

(Procedure Division)

```

      ADD S, N GIVING RMOD.
RCV.   CALL "MRCV" USING STAT, RETBUF, RMOD, Z-VAL, Z-VAL,
      .   BUFFER, Z-VAL, Z-VAL, SLPTIME.
      IF STAT = 5 GO TO RCV.
      MOVE N TO RMOD.
      CALL "MRCV" USING STAT, RETBUF, RMOD, Z-VAL, Z-VAL,
      .   BUFFER, BUFLen, Z-VAL, Z-VAL.
      IF STAT NOT EQUAL 0 GO TO ERR-RTN.
      MOVE RETBUF-CHAR(3) TO ABC.
      CALL "BYTE-VAL" USING ABC, DEF.
      ADD DEF, 10 GIVING JMP-VAL.
      GO TO LAPROC, DIPROC, LSPROC, INPROC, DMPROC,
      .   CRPROC, CCPROC, CIPROC, LMPROC
      .   DEPENDING ON JMP-VAL.
LAPROC. (processing for Link Abort Message)
DIPROC. (processing for Disconnect Message)
      .
      .
      .

```

7.7.1 Format of Received Local Data Message

The format of the information returned to the *retbuf* buffer for a Local Data Message is shown below. Up to 512 bytes of user data can be delivered to the data buffer specified in the Receive call (MRCV).

| Bytes | Data Type | Description |
|-------|-----------------|---|
| 1-2 | — | Not meaningful - should be ignored. |
| 3 | one-byte binary | The value -1, the function code for a Local Data Message. |
| 4 | one-byte binary | Two times the job number of the local sender. |
| 5 | one-byte binary | The programmer number of the local sender (from the project-programmer number). |
| 6 | one-byte binary | The project number of the local sender (from the project-programmer number). |
| 7-8 | — | Not meaningful - should be ignored. |
| 9-10 | one-word COMP | The number of bytes of user data not delivered to the buffer. This data has been either discarded or saved for a later Receive call, depending on how the <i>t</i> flag was set in the <i>rmod</i> argument in the MRCV call. |
| 11-12 | — | Not meaningful - should be ignored. |
| 13-14 | one-word COMP | The number of bytes of user data transferred to the buffer on this Receive call. |
| 15-20 | — | Not meaningful - should be ignored. |
| 21-40 | (whatever sent) | The data passed as parameters by the sender of this message. The system pads any unused bytes with binary zeros to a length of 20 bytes. |

7.7.2 Format of Received Connect Initiate Message

The format of the information returned to the *retbuf* buffer for a received Connect Initiate Message is shown below. A 120-byte Connect Data Block and up to 16 bytes of user data are also returned to the data buffer specified in the Receive call. The format of a received Connect Data Block is shown in Figure 7-2 and described in detail in Table 7-5.

| Bytes | Data Type | Description |
|-------|-----------------|---|
| 1-2 | — | Not meaningful - should be ignored. |
| 3 | one-byte binary | The value -2, the function code for a Connect Initiate Message. |
| 4 | — | Not meaningful - should be ignored. |
| 5-6 | one-word COMP | Local link address. Identifies the link for the local program's subsequent Connect Confirm or Connect Reject Message. |
| 7-8 | — | Not meaningful - should be ignored. |
| 9-10 | one-word COMP | The number of bytes of user data not delivered to the buffer. This data has been either discarded or saved for a later Receive call, depending on how the <i>t</i> bit was set in the <i>rmod</i> argument in the MRCV call. |
| 11-12 | — | Not meaningful - should be ignored. |
| 13-14 | one-word COMP | The number of bytes of user data transferred to the buffer on this Receive call. |
| 15-22 | — | Not meaningful - should be ignored. |
| 23 | one-byte binary | Remote link modifier. Indicates the flow control in effect for the remote program. If = 0, no flow control = 1, segment flow control = 2, message flow control |
| 24 | — | Not meaningful - should be ignored. |
| 25-26 | one-word COMP | Receive maximum (in bytes). The maximum amount of user data that will be transmitted to the local program in one Network Data Message over this link. This value is supplied by the local NSP and is determined by the size of the receive buffers that it allocates. The maximum size of the receive buffers is set for NSP by Network Management. The local program can further limit the amount of user data to be received by specifying a smaller <i>rmax</i> in the Connect Confirm Message that accepts this request for a logical link (see Section 7.6.3). |
| 27-28 | one-word COMP | Transmit maximum, in bytes. The maximum amount of user data that can be transmitted by the program in one Network Data Message over this link. This limit has been imposed by either the remote program or the remote NSP. The local NSP enforces this limit for all Network Data Messages sent by the local program on this link. |
| 29-40 | — | Not meaningful - should be ignored. |

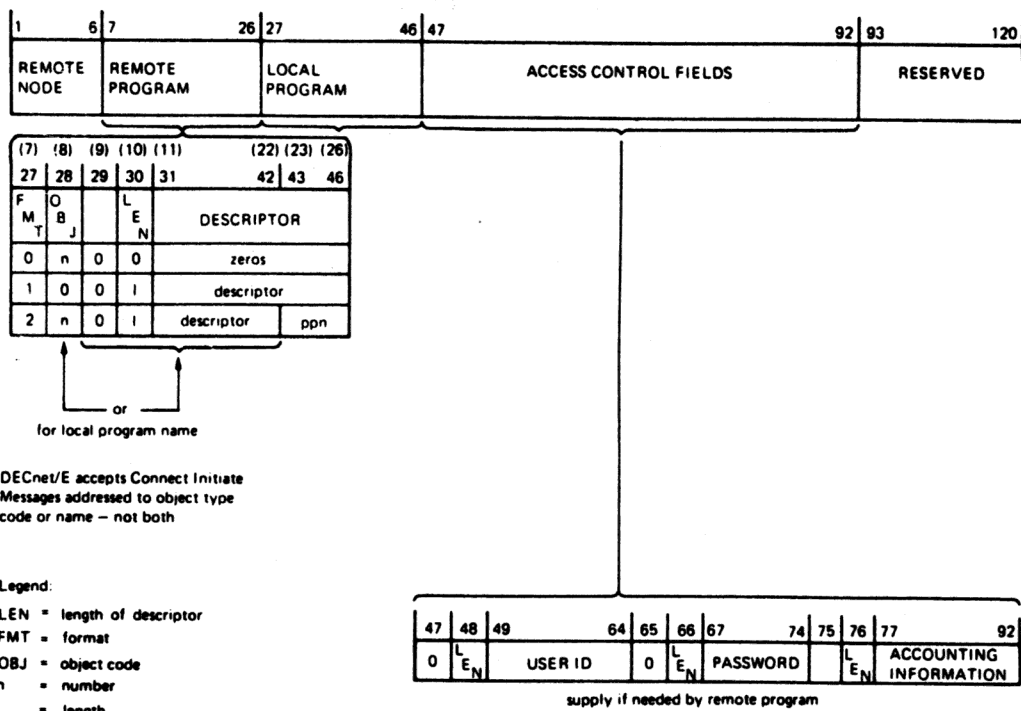


Figure 7-2: Format of Received Connect Data Block

Table 7-5: Format of Received Connect Data Block

| Bytes | Content | | | | | | | | |
|-------|---|-------|---------|---|---|---|--|------|-------------------------------------|
| 1-6 | Remote Node. One to six uppercase, alphanumeric ASCII characters naming the remote node from which the Connect Initiate Message was received. The name will contain at least one alphabetic character. Node names shorter than six characters will be left-justified and padded to six characters with spaces. | | | | | | | | |
| 7-26 | Remote Program. Identifies the remote program that sent the Connect Initiate Message as either an object type (general function) or a specific entity (program, task, or process). This identification is given in one of three formats: <div> <div> Format 0 <p>The remote program is identified by object type code alone. The fields are:</p> <table> <tr> <th>Bytes</th><th>Content</th></tr> <tr> <td>7</td><td>One-byte binary value of zero, to indicate a format 0 name.</td></tr> <tr> <td>8</td><td>One-byte binary value giving the object type code of the remote program.</td></tr> <tr> <td>9-26</td><td>Not meaningful - should be ignored.</td></tr> </table> </div> <div> Format 1 <p>The remote program is identified by a descriptor. The fields are:</p> </div> </div> | Bytes | Content | 7 | One-byte binary value of zero, to indicate a format 0 name. | 8 | One-byte binary value giving the object type code of the remote program. | 9-26 | Not meaningful - should be ignored. |
| Bytes | Content | | | | | | | | |
| 7 | One-byte binary value of zero, to indicate a format 0 name. | | | | | | | | |
| 8 | One-byte binary value giving the object type code of the remote program. | | | | | | | | |
| 9-26 | Not meaningful - should be ignored. | | | | | | | | |

(continued on next page)

Table 7-5 (Cont.): Format of Received Connect Data Block

Bytes Content

| | |
|-------|--|
| 7 | One-byte binary value of one, to indicate a format 1 name. |
| 8 | Normally zero (one-byte binary value). This byte will usually be zero unless the remote DECnet system allows identification by both name and object type code on outgoing connect requests (DECnet/E does not). If nonzero, this byte is interpreted as the object type code of the remote program that sent the Connect Initiate Message. |
| 9 | Not meaningful - should be ignored. |
| 10 | One-byte binary value giving the length of the descriptor, 0 to 16 bytes. Gives the number of characters to be interpreted as the program descriptor in the following field (bytes 11-26). |
| 11-26 | Remote program descriptor (left-justified). The descriptor used to register the remote program for network operations. For remote DECnet/E systems, this is the logical name declared in the Declare Receiver call. |

Format 2

This format identifies the remote program by name or object type code and by the group and user codes under which the remote program is running on the remote system. (The group and user codes refer to what is called the project-programmer number on some DIGITAL systems, including RSTS/E. Other systems refer to these codes as the user identification code, or UIC.) DECnet/E systems use this format to identify the sender of a Connect Initiate Message.

The subfields are defined as follows:

Bytes Content

| | |
|-------|---|
| 7 | One-byte binary value of two, to indicate a format 2 name. |
| 8 | One-byte binary value giving the object type code of the remote program, if applicable. If zero, the descriptor alone identifies the remote program. |
| 9 | Not meaningful - should be ignored. |
| 10 | One-byte binary value giving the length of the descriptor, 0 to 12 bytes. Gives the number of characters to be interpreted as the program descriptor in the following field (bytes 11-22). |
| 11-22 | Remote program descriptor, if this is used instead of object type code. The descriptor will be left-justified in the field. |
| 23-24 | Remote group code. One-word binary value giving the group number under which the remote program is running. If the remote system is a DECnet/E system, byte 23 is the project number in a project-programmer number. |
| 25-26 | Remote user code. One-word binary value giving the user number under which the remote program is running. If the remote system is a DECnet/E system, byte 25 is the programmer number in a project-programmer number. |

27-46 Local Program. This field indicates how the remote program addressed the local program in the Connect Initiate Message. DECnet/E allows the local program to be addressed by object type code, or by name, but not both. Again, three formats can be used. The subfields are as follows:

(continued on next page)

Table 7-5 (Cont.): Format of Received Connect Data Block

| Bytes | Content |
|---------------|---|
| 27 | Format type (one-byte binary value). Equal to 0, 1, or 2, depending on the format used by the remote program to address the local program. |
| 28 | Object type code, if used by the remote program to address the local program (formats 0 or 2). One-byte binary value. |
| 29 | Always zero. |
| 30 | One-byte binary value giving the length of name, in bytes. Will be 0 for format 0, 1 to 6 for format 1, 0 to 6 for format 2. (DECnet/E program names are limited to a maximum of six characters by the Declare Receiver call.) |
| 31-36 | Local program name, if used by the remote program to address the local program (formats 1 or 2). |
| 37-42 | Always zero. |
| 43-44 | Local group code. One-byte binary value in byte 43 giving the local program's project number, if used by the remote program to address the local program (format 2). Zero otherwise (formats 0 or 1). Not used by NSP to identify the local program, but passed on here. |
| 45-46 | Local user code. One-byte binary value in byte 45 giving the local program's programmer number, if used by the remote program to address the local program (format 2). Zero otherwise (formats 0 or 1). Not used by NSP to identify the local program but passed on here. |
| 47-92 | Access Control Fields. These bytes can be used to define the remote program's access rights to the local program's services. The idea is analogous to logging in to the local system. The most rigorous checking done in a normal log-in on DIGITAL systems involve user identification, password, and account number, so three fields are defined here for those items. However, overall DECnet design does not specifically require use of these fields at all nor does it require that the fields be used as described. A DECnet/E system simply passes this information on to the receiving program. |
| Bytes | Content |
| 47 | Not meaningful - should be ignored. |
| 48 | One-byte binary value giving the length of the user identification information specified in the next field, 0 to 16 bytes. |
| 49-64 | User Identification. This field identifies the person or program requesting the logical link. In RSTS/E terms, the user ID is analogous to the project-programmer number used at log-in. |
| 65 | Not meaningful - should be ignored. |
| 66 | One-byte binary value giving the length of the password information specified in the next field, 0 to 8 bytes. |
| 67-74 | Password. A password is often used to verify access to a system under a particular user identification. In RSTS/E, this password is analogous to the password used in conjunction with the project-programmer number at log-in. |
| 75 | Not meaningful - should be ignored. |
| 76 | One-byte binary value giving the length of the accounting information specified in the following field, 0 to 16 bytes. |
| 77-92 | Accounting information. Many systems require a billing account number in addition to user identification and password. The accounting field is provided for this purpose. |
| 93-120 | Reserved for future use. Should be ignored. |

7.7.3 Format of Received Connect Confirm Message

A received Connect Confirm Message is a remote program's acceptance of a Connect Initiate Message sent by the receiving program. The link is identified by the user link address specified by the receiving program in the Connect Initiate Message. The format of the information returned to the *retbuf* buffer is described below. Up to 16 bytes of user data can also be returned to the data buffer specified in the Receive call.

| Bytes | Data Type | Description |
|-------|-----------------|---|
| 1-2 | — | Not meaningful - should be ignored. |
| 3 | one-byte binary | The value -3, the function code for a Connect Confirm Message. |
| 4 | one-byte binary | User link address. |
| 5-8 | — | Not meaningful - should be ignored. |
| 9-10 | one-word COMP | The number of bytes of user data not delivered to the data buffer. This data has been either discarded or saved for a later Receive call, depending on how the <i>t</i> bit was set in the <i>rmod</i> argument in the MRCV call. |
| 11-12 | — | Not meaningful - should be ignored. |
| 13-14 | one-word COMP | The number of bytes of user data transferred to the data buffer on this Receive call. |
| 15-22 | — | Not meaningful - should be ignored. |
| 23 | one-byte binary | Remote link modifier. Indicates the flow control in effect for the remote program. If = 0, no flow control = 1, segment flow control = 2, message flow control |
| 24 | — | Not meaningful - should be ignored. |
| 25-26 | one-word COMP | Receive maximum (in bytes). The maximum amount of user data that will be transmitted to the local program in one Network Data Message over this link. This will be (1) equal to the <i>rmax</i> value specified in the Connect Initiate Message for this link, or (2) a lesser value supplied by the local NSP, determined from the size of the receive buffers it allocates. |
| 27-28 | one-word COMP | Transmit maximum (in bytes). The maximum amount of user data that can be transmitted by the local program in one Network Data Message over this link. This limit has been imposed by either the remote program or the remote NSP. The local NSP enforces this limit for all Network Data Messages sent by the local program on this link. |
| 29-40 | — | Not meaningful - should be ignored. |

7.7.4 Format of Received Connect Reject Message

A received Connect Reject Message is a rejection of the receiving program's request for a logical link. It can be from the remote program or from the local or remote NSP. The link is identified by the user link address previously established in the receiving program's Connect Initiate Message. The format of the information returned to the *retbuf* buffer is described below. Rejections from NSP will contain a reason code explaining the reason for the rejection and will not contain user data. Rejections from the remote program can contain up to 16 bytes of user data, delivered to the data buffer defined in the Receive call.

| Bytes | Data Type | Description |
|-------|-----------------|---|
| 1-2 | — | Not meaningful - should be ignored. |
| 3 | one-byte binary | The value -4, the function code for a Connect Reject Message. |
| 4 | one-byte binary | User link address, previously defined in the local program's Connect Initiate Message. |
| 5-8 | — | Not meaningful - should be ignored. |
| 9-10 | one-word COMP | The number of bytes of user data not delivered to the data buffer. This data has been either discarded or saved for a later Receive call, depending on how the <i>t</i> bit was set in the <i>rmod</i> argument in the MRCV call. |
| 11-12 | — | Not meaningful - should be ignored. |
| 13-14 | one-word COMP | The number of bytes of user data transferred to the data buffer on this Receive call. |
| 15-22 | — | Not meaningful - should be ignored. |
| 23-24 | one-word COMP | Identifies the reason for the rejection (see Appendix B). |
| 25-40 | — | Not meaningful - should be ignored. |

7.7.5 Format of Received Network Data Message

A received Network Data Message contains user data from a remote program. The format of the information returned to the *retbuf* buffer described below. The user data accompanying the Network Data Message is delivered to the data buffer specified in the Receive call.

| Bytes | Data Type | Description |
|-------|-----------------|--|
| 1-2 | — | Not meaningful - should be ignored. |
| 3 | one-byte binary | The value -5, the function code for a Network Data Message. |
| 4 | one-byte binary | User link address, previously defined in the local program's Connect Initiate or Connect Confirm Message. |
| 5-8 | — | Not meaningful - should be ignored. |
| 9-10 | one-word COMP | The number of bytes of user data not delivered to the data buffer. This data has been either discarded or saved for a later Receive call, depending on how the <i>t</i> bit was set in the <i>rmod</i> argument in the MRCV call. |
| 11-12 | — | Not meaningful - should be ignored. |
| 13-14 | one-word COMP | The number of bytes of user data transferred to the data buffer on this Receive call. |
| 15-20 | — | Not meaningful - should be ignored. |
| 21 | one-byte binary | Data message flags. Low-order bits indicate whether this is the beginning, middle, end, or sole segment of a logical message. If = 0, Middle segment of message = 1, First segment of message = 2, Last segment of message = 3, Sole segment of message DECnet/E uses bit 1 (set if this byte = 2 or 3) to determine the end of a logical message if the local program requested message flow control when the logical link was established. It does not use bit 0 but passes it on to the receiving program. |
| 22-40 | — | Not meaningful - should be ignored. |

7.7.6 Format of Received Interrupt Message

The local NSP places Interrupt Messages from a remote program at the head of the pending message queue, behind the first message and behind any other pending Interrupt Messages. The format of the data returned to the *retbuf* buffer is described below. Up to 16 bytes of user data can be delivered to the data buffer specified in the Receive call.

| Bytes | Data Type | Description |
|-------|-----------------|--|
| 1-2 | — | Not meaningful - should be ignored. |
| 3 | one-byte binary | The value -6, the function code for an Interrupt Message. |
| 4 | one-byte binary | User link address, previously defined in the local program's Connect Initiate or Connect Confirm Message. |
| 5-8 | — | Not meaningful - should be ignored. |
| 9-10 | one-word COMP | The number of bytes of user data not delivered to the buffer. This data has been either discarded or saved for a later Receive call, depending on how the <i>t</i> bit was set in the <i>rmod</i> argument in the MRCV call. |
| 11-12 | — | Not meaningful - should be ignored. |
| 13-14 | one-word COMP | The number of bytes of user data transferred to the buffer on this Receive Call. |
| 15-40 | — | Not meaningful - should be ignored. |

7.7.7 Format of Received Link Service Message

A Link Service Message provides status information for a particular logical link. A Link Service Message is only returned on a Receive call when the pending message queue is empty after a condition has cleared that previously inhibited the receiving program from sending (see Section 6.2). Status information with the same format as a Link Service Message can also be returned to the *retbuf* buffer on successful completion of a Send Link Service, Send Network Data, or Send Interrupt call. The format of the information returned is described below. No data is delivered to the data buffer with a Link Service Message.

| Bytes | Data Type | Description | | | | | | | | | | | | | | | |
|---|-----------------|---|-------|-------|--|---|---|-----------------------------------|---|---|---|---|---|---|---|---|--|
| 1-2 | — | Not meaningful - should be ignored. | | | | | | | | | | | | | | | |
| 3 | one-byte binary | The value -7, the function code for a Link Service Message. | | | | | | | | | | | | | | | |
| 4 | one-byte binary | User link address, previously defined in the local program's Connect Initiate or Connect Confirm Message. | | | | | | | | | | | | | | | |
| 5-8 | — | Not meaningful - should be ignored. | | | | | | | | | | | | | | | |
| 9-10 | one-word COMP | Always zero for Link Service. (These bytes are the remainder field in other returned messages.) | | | | | | | | | | | | | | | |
| 11-12 | — | Not meaningful - should be ignored. | | | | | | | | | | | | | | | |
| 13-14 | one-word COMP | Always zero for Link Service. (These bytes are the length field in other returned messages.) | | | | | | | | | | | | | | | |
| 15-20 | — | Not meaningful - should be ignored. | | | | | | | | | | | | | | | |
| 21 | one-byte binary | Local status flags. (See Section 7.1 for tips on testing bits.) | | | | | | | | | | | | | | | |
| If bit 0 = 1, outgoing Network Data Messages are inhibited. (A Link Service Message will be delivered when the condition clears.) | | | | | | | | | | | | | | | | | |
| If bit 1 = 1, outgoing Interrupt and Link Service Messages are inhibited. (A Link Service Message will be delivered when the condition clears.) | | | | | | | | | | | | | | | | | |
| Bits 6 and 7 indicate the local NSP's backpressure status: | | | | | | | | | | | | | | | | | |
| <table> <tr> <th>bit 7</th><th>bit 6</th><th></th></tr> <tr> <td>0</td><td>0</td><td>Incoming Data Message flow is on.</td></tr> <tr> <td>0</td><td>1</td><td>Incoming Data Message flow is on but the local NSP will turn this link off if another Data Message is received before the local program's receive queue is emptied.</td></tr> <tr> <td>1</td><td>0</td><td>Incoming Data Message flow has been turned off by the local NSP. The local program's receive queue must be emptied before the link will be turned on.</td></tr> <tr> <td>1</td><td>1</td><td>Incoming Data Message flow has been turned off by the local NSP but flow is scheduled to turn on as soon as system buffers become available.</td></tr> </table> | | | bit 7 | bit 6 | | 0 | 0 | Incoming Data Message flow is on. | 0 | 1 | Incoming Data Message flow is on but the local NSP will turn this link off if another Data Message is received before the local program's receive queue is emptied. | 1 | 0 | Incoming Data Message flow has been turned off by the local NSP. The local program's receive queue must be emptied before the link will be turned on. | 1 | 1 | Incoming Data Message flow has been turned off by the local NSP but flow is scheduled to turn on as soon as system buffers become available. |
| bit 7 | bit 6 | | | | | | | | | | | | | | | | |
| 0 | 0 | Incoming Data Message flow is on. | | | | | | | | | | | | | | | |
| 0 | 1 | Incoming Data Message flow is on but the local NSP will turn this link off if another Data Message is received before the local program's receive queue is emptied. | | | | | | | | | | | | | | | |
| 1 | 0 | Incoming Data Message flow has been turned off by the local NSP. The local program's receive queue must be emptied before the link will be turned on. | | | | | | | | | | | | | | | |
| 1 | 1 | Incoming Data Message flow has been turned off by the local NSP but flow is scheduled to turn on as soon as system buffers become available. | | | | | | | | | | | | | | | |

(continued on next page)

| Bytes | Data Type | Description |
|-------|-----------------|--|
| 22 | one-byte binary | Remote status flags. Only one bit is currently defined for this field. If bit 7 = 128, the link has been turned off by the remote system due to a backpressure condition. No Data Messages can be sent until the remote system turns the link back on. |
| 23 | one-byte binary | Local data request count. The count of segments or logical messages currently requested by the local program, reflecting the actual number of segments or logical messages requested but not yet received. |
| 24 | one-byte binary | Local interrupt request count. The count of Interrupt Messages currently requested by the local program, reflecting the actual number of interrupts requested but not yet received. This count will never exceed one. |
| 25 | one-byte binary | Remote data request count. The count of segments or logical messages currently requested by the remote program, reflecting the actual count of segments or messages requested but not yet sent. |
| 26 | one-byte binary | Remote interrupt request count. The count of Interrupt Messages currently requested by the remote end of the logical link, reflecting the actual count of interrupts requested but not yet sent. |
| 27 | one-byte binary | Data transmit queue maximum. The maximum number of data messages that can be queued waiting for acknowledgment from the remote system. This is a constant set by the system manager when NSP is enabled. It applies individually to all active logical links. |
| 28 | one-byte binary | Data transmit queue count. The number of data messages currently in the data transmit queue for this logical link. It is a count of data messages waiting for acknowledgment from the remote NSP. |
| 29 | one-byte binary | Interrupt/Link Service transmit queue maximum. The maximum number of Interrupt or Link Service Messages that can be queued waiting for acknowledgment from the remote system. This is a constant set by the system manager when NSP is enabled. It applies individually to all active logical links. |
| 30 | one-byte binary | Interrupt/Link Service transmit queue count. The number of Interrupt and Link Service Messages currently in the Interrupt/Link Service transmit queue for this logical link. It is a count of Interrupt and Link Service Messages waiting for acknowledgment from the remote NSP. |
| 31 | one-byte binary | Message maximum. The maximum number of incoming messages that will be queued for the local program. This maximum is set by the program in its Declare Receiver call. |
| 32 | one-byte binary | Message count. The total number of messages currently in the pending message queue for the local program. |
| 33 | one-byte binary | Message count for this logical link. The number of messages currently in the pending message queue for the logical link identified by the <i>ula</i> in byte 4. |
| 34-40 | — | Not meaningful – should be ignored. |

7.7.8 Format of Received Disconnect Message

A received Disconnect Message indicates that an established logical link with a remote program has been disconnected by the remote program. All other messages sent by the remote program over this logical link have been received by the local program. The format of the data returned to the *retbuf* buffer is defined below. Up to 16 bytes of user message data can be returned to the data buffer specified in the Receive call.

| Bytes | Data Type | Description |
|-------|-----------------|---|
| 1-2 | — | Not meaningful - should be ignored. |
| 3 | one-byte binary | The value -8, the function code for a Disconnect Message. |
| 4 | one-byte binary | User link address, previously defined in the local program's Connect Initiate or Connect Confirm Message. |
| 5-8 | — | Not meaningful - should be ignored. |
| 9-10 | one-word COMP | The number of bytes of user data not delivered to the data buffer. This data has been either discarded or saved for a later Receive call, depending on how the <i>t</i> bit was set in the <i>rmod</i> argument in the MRCV call. |
| 112 | — | Not meaningful - should be ignored. |
| 13-14 | one-word COMP | The number of bytes of user data transferred to the data buffer on this Receive call. |
| 15-40 | — | Not meaningful - should be ignored. |

7.7.9 Format of Received Link Abort Message

A received Link Abort Message indicates that a logical link has been terminated by the remote program or by the local or remote NSP. The format of the data returned to the *retbuf* buffer is defined below. Up to 16 bytes of user message data can be returned to the data buffer specified in the Receive call.

| Bytes | Data Type | Description |
|-------|-----------------|--|
| 1-2 | — | Not meaningful - should be ignored. |
| 3 | one-byte binary | The value -9, the function code for a Link Abort Message. |
| 4 | one-byte binary | User link address (if any), previously defined in the local program's Connect Initiate or Connect Confirm Message. |
| 5-6 | one-word COMP | Local link address, established by the local NSP for this link. It identifies the link if the remote program sent a Connect Initiate Message but the local program has not yet responded with a Connect Confirm or Connect Reject Message. No <i>ula</i> exists under these conditions to identify the link being aborted. |
| 7-8 | — | Not meaningful - should be ignored. |
| 9-10 | one-word COMP | The number of bytes of user data not delivered to the data buffer. This data has been either discarded or saved for a later Receive call, depending on how the <i>t</i> flag was set in the <i>rmod</i> argument in the MRCV call. |
| 11-12 | — | Not meaningful - should be ignored. |
| 13-14 | one-word COMP | The number of bytes of user data transferred to the data buffer on this Receive call. |
| 15-22 | — | Not meaningful - should be ignored. |
| 23-24 | one-word COMP | Abort reason. If nonzero, the link was aborted by either the local or remote NSP and the reason is specified in these two bytes. The reasons that apply for a link abort are listed in Appendix B. If reason is zero, the link was aborted at the request of the remote program. |
| 25-40 | — | Not meaningful - should be ignored. |

Appendix A

Object Type Codes

| Code | Type of Process |
|-------------|--|
| 000 | General Task, User Process |
| 001 | File Access (FAL/DAP Version 1) |
| 002 | Unit Record Services (URD) |
| 003 | Application Terminal Services (ATS) |
| 004 | Command Terminal Services (CTS) |
| 005 | RSX-11M Task Control, Version 1 |
| 006 | Operator Services Interface |
| 007 | Node Resource Manager |
| 008 | IBM 3270-BSC Gateway |
| 009 | IBM 2780-BSC Gateway |
| 010 | IBM 3790-SDLC Gateway |
| 011 | TPS Application |
| 012 | RT-11 DIBOL Application |
| 013 | TOPS-20 Terminal Handler |
| 014 | TOPS-20 Remote Spooler |
| 015 | RSX-11M Task Control, Version 2 |
| 016 | TLK Utility (LSN on DECnet/E) |
| 017 | File Access (FAL/DAP Version 4 and later) |
| 018 | RSX-11S Remote Task Loader |
| 019 | Network Management Listener (NICE Process) |
| 020 | RSTS/E Media Transfer Program (NETCPY) |
| 021 | Reserved for DECnet use |
| 022 | Mail Listener |
| 023 | Host Terminal Handler (NPKDVR) |
| 024 | Concentrator Terminal Handler |

(continued on next page)

| | |
|---------|----------------------------------|
| 025 | Loopback Mirror |
| 026 | Event Receiver |
| 027 | VAX/VMS Personal Message Utility |
| 028 | File Transfer Spooler (FTS) |
| 029-062 | Reserved for DECnet use |
| 063 | DECnet test tool (DTR) |
| 064-127 | Reserved for DECnet use |
| 128-255 | Reserved for customer extensions |

Appendix B

NSP Reason Codes for Connect Reject and Link Abort

| Code | Reason |
|-------------|--|
| 000 | No error — user-initiated reject or abort |
| 001 | Resource allocation failure |
| 002 | Destination node does not exist |
| 003 | Node shutting down |
| 004 | Destination program does not exist |
| 005 | Invalid destination name or source name field |
| 006 | Destination program's queue full |
| 007 | Unspecified error condition |
| 008 | Third party aborted logical link |
| 009 | User-initiated link abort |
| 010-020 | Reserved |
| 021 | Invalid destination address in Connect Initiate Message |
| 022 | Invalid destination address in Connect Confirm Message |
| 023 | Source address zero in Connect Initiate or Connect Confirm Message |
| 024 | Flow control violation — invalid request count in Link Service Message |
| 025-031 | Reserved |
| 032 | Too many connects to node |
| 033 | Too many connects to destination program |
| 034 | Access not permitted |
| 035 | Logical link services mismatch |
| 036 | Invalid accounting information |
| 037 | Segment size too small |
| 038 | User aborted, timed out, or canceled link |

- 039 No path to node
- 040 Flow control failure — data received when request count zero
- 041 No current link (cannot recognize destination address)
- 042 Confirmation from remote system of Disconnect Message
- 043 Image data field too long

INDEX

A

Access,
 consistency with receiver identity, 7-8
 specifying, 7-7
 types permitted, 7-7
Access control fields, 7-28, 7-85
Accounting information, 7-29, 7-35, 7-68
Adjacent node, 1-1, 1-4
Automatic startup, 2-3, 5-6, 6-3, 6-4
 ordering of objects, 6-5
 starting line number, 6-4
 timing, 6-5

B

Backpressure flow control, 3-9, 6-13
 clearing, 6-13
 effect on DECnet/E programs, 6-13
 use of Link Service Message, 6-13
Buffer,
 declared maximum, 7-9
 limited space, 7-24, 7-32, 7-60
 monitor pool, 4-7, 5-2, 5-4, 5-5, 7-9
 NSP maximum, 5-7

C

CALL statement, 7-1
CNTACP call,
 See Concise interface
CNTCON call,
 See Concise interface
CNTDIS call,
 See Concise interface
CNTRCV call,
 See Concise interface
CNTSND call,
 See Concise interface
Comparison of interfaces, 2-7
Compiling, 3-2, 7-5
Concise interface, 2-3
 accounting information, 3-4, 3-7
 comparison with extended interface, 2-7
 Connect Accept (CNTACP), 2-4, 3-6
 Connect Request (CNTCON), 2-4, 3-3
 deadlock situations, 2-4
 Disconnect (CNTDIS), 2-4, 3-13
 files, 2-3
 form of calls, 3-1
 password, 3-4, 3-7
 PPN, 3-4, 3-7

Receive Record (CNTRCV), 2-4, 3-11
 receiver identity, 3-3, 3-6
 record length, 2-5
 records, 2-3
Send Record (CNTSND), 2-4, 3-9
 synchronous I/O, 2-4
Connect Confirm Message, 2-6, 5-8
 received format, 7-69
 send call, 7-32
Connect Data Block, 5-7, 5-8, 6-6, 7-25, 7-26, 7-66
Connect Initiate Message, 2-6, 5-7
 local program identification, 7-28, 7-67
 received format, 7-65
 remote program identification, 7-26, 7-67
 send call, 7-24
Connect Reject Message, 2-6, 5-9
 received format, 7-70
 send call, 7-36
CORCMN, 6-5
Core common, 6-5
Cost, 1-5
CRC, 1-4
CTRL/C, 5-5
Cyclic Redundancy Check,
 See CRC

D

DAP, 1-4
Data Access Protocol,
 See DAP
Data Division, 3-2, 7-2
Data transmit queue, 6-10, 7-53, 7-74
DDCMP, 1-3, 1-5
Deadlock situations, 2-4
Declare Receiver, 2-5, 5-2, 5-3
 access, 7-7
 declaring intended usage, 5-4
 link maximum, 5-4
 message maximum, 5-4
 receiver identity,
 See Receiver identity
 sleep bit, 7-8
 system call, 7-6
Default project-programmer number, 5-5, 6-6, 7-14
DEFINE OBJECT, 6-4
Device controllers, 1-5
DIGITAL Data Communications Message Protocol,
 See DDCMP
DIGITAL Network Architecture,
 See DNA
Disconnect Message, 2-6, 5-11
 affected by transmit queues, 6-11
 received format, 7-75
 send call, 7-21

DMx controllers, 1-5
DNA, 1-2, 1-3, 5-6, 7-17

E

End node, 1-4
ERRCPY, 7-59
Errors,
 Connect Accept (Concise interface), 3-7, 3-8
 Connect Request (Concise interface), 3-5
 Declare Receiver, 7-10
 Disconnect (Concise interface), 3-13
 Get Local Node Parameters, 7-14
 Log User Event, 7-17
 Receive, 7-58
 Receive Record (Concise interface), 3-12
 Remove Receiver, 7-12
 Send Connect Confirm Message, 7-32
 Send Connect Initiate Message, 7-24
 Send Connect Reject Message, 7-36
 Send Disconnect Message, 7-52
 Send Interrupt Message, 7-44
 Send Link Abort Message, 7-55
 Send Link Service Message, 7-48
 Send Local Data Message, 7-21
 Send Network Data Message, 7-39
 Send Record (Concise interface), 3-9
Events,
 See User events
Extended interface, 2-5
 comparison with concise interface 2-7

F

File Access Listener,
 See NFT/FAL
Flow control, 4-7, 5-10, 6-9
 backpressure,
 See Backpressure flow control
 Interrupt Message,
 See Interrupt Message flow control
 local congestion, 6-21
 local modifier, 7-24, 7-32
 program regulated,
 See Program-regulated flow control
 programming hints, 6-20
 reasons for inhibiting transmission, 6-20
 remote congestion, 6-21
 remote modifier, 7-27, 7-66
 transmit queue management,
 See Transmit queue management
 use of Link Service Message, 6-21
Full duplex transmission, 4-2

H

Hop, 1-2

Interrupt Message, 2-6, 4-7, 5-10
 queuing, 4-7, 5-11, 7-45, 7-72
 received format, 7-72
 reenabling, 5-11, 6-19
 request count, 7-45, 7-72
 send call, 7-21
Interrupt Message flow control, 5-10, 6-19
 effect on DECnet/E programs, 6-19
 use of Link Service Message, 6-19
Interrupt/Link Service queue, 6-10, 7-45
Interrupts,
 See Interrupt Message

J

Job,
 determining number, 6-2
 killing, 6-5
 number in Receive, 7-12
 number of local receiver, 5-6, 7-21
 privilege level, 6-5
 Receiver ID Block, 5-5, 7-6
 receiving from system job 0,
 Remove Receiver, 5-5, 7-12

K

Killing a job, 6-5

L

Link,
 logical, 1-1, 1-3, 1-5
 physical, 1-1, 1-3, 1-5
Link Abort Message, 2-7, 5-12
 received format, 7-76
 send call, 7-55
 unaffected by transmit queues, 6-11
Link master, 4-1
Link Service Message, 2-6, 5-10, 5-11
 purpose, 5-11, 6-13, 6-15, 6-17, 6-19
 received format, 7-73
 send call, 7-48
Link status bit, 7-7, 7-8
Link status flags,
 local, 7-73
 remote, 7-74
Link status information, 5-10, 5-11, 7-49
Linking, 3-2, 7-5
LLA, 4-2, 5-7, 5-8, 5-9, 5-12
Local communication, 5-1, 5-6, 7-21

Local Data Message, 2-6, 5-6
 received format, 7-64
 send call, 7-21
 Local link address,
 See LLA
 Local link status flags, 7-73
 Local node, 1-1, 5-5
 Local node parameters, 2-5, 5-2, 5-5
 system call, 7-14
 use of, 6-6
 Local program, 5-3, 5-4
 identification, 7-28
 Log-in information, 5-6, 6-6
 Logical full duplex transmission, 4-2
 Logical link, 2-1, 4-1
 aborting, 5-5
 See Link Abort Message
 accepting,
 See Connect Confirm Message
 active, 5-5
 creating, 2-2, 4-1
 declared maximum,
 declared message maximum, 6-2, 6-3, 7-8, 7-9
 DECnet/E limitations, 4-9
 disconnecting,
 See Disconnect Message
 efficient line usage, 4-4
 identifiers, 4-2, 4-4
 message count, 7-74
 message streams, 4-5, 5-10
 receive maximum,
 See Receive maximum
 rejecting,
 See Connect Reject Message
 requesting,
 See Connect Initiate Message
 transmit maximum,
 See Transmit maximum
 Logical message, 6-15
 end of message segment, 6-16
 request count, 6-16

M

MACRO-11 subprograms, 1-4, 2-5, 3-1, 3-2, 6-2,
 6-5, 6-6, 7-1, 7-2
 Master-slave communication, 5-12
 MDCL call,
 See Declare Receiver
 Message,
 count for logical link, 7-74
 count of pending messages, 7-74
 declared maximum, 5-4, 6-2, 6-4, 7-9, 7-17
 limit of transmit queues, 6-11, 7-69
 logical,
 See Logical message
 request count, 6-16
 Message flow control, 6-15

Message streams, 2-1, 4-5, 4-6, 5-10
 MOVE statement, 3-2, 7-2
 MRCV call,
 See Receive
 MREM call,
 See Local Data Message
 Multiple copies of objects, 2-3, 5-4, 6-2, 6-3
 under BATCH, 6-4
 via automatic startup, 6-4
 via terminal users, 6-4
 Multipoint lines, 1-5

N

NCP, 1-4, 1-5, 6-4, 6-10
 NET, 1-4
 NETCPY, 1-4, 6-7
 Network, 1-1, 1-2
 as a file, 2-3
 Network addressing, 2-3, 5-4, 6-1
 declaring identity,
 See Receiver identity
 handling incoming requests, 6-2
 permitted by DECnet/E, 6-2
 to name, 6-3
 to object type code, 6-2
 Network Command Terminal Utility,
 See NET
 Network Communication Utility,
 See TLK/LSN
 Network Control Program,
 See NCP
 Network Copy Utility,
 See NETCPY
 Network Data Message, 2-6, 5-10
 received format, 7-71
 send call, 7-39
 Network diameter, 1-2
 Network File Transfer,
 See NFT/FAL
 Network Information and Control Exchange Protocol,
 See NICE
 Network program, 5-3, 5-4
 Network services Protocol,
 See NSP
 NFT/FAL, 1-4, 6-6, 7-8, 7-36
 NICE, 7-13, 7-17
 Node, 1-1
 address, 5-5, 7-14
 adjacent, 1-1, 1-4
 alias, 5-5, 7-14
 local, 1-1, 5-5
 name, 3-3, 5-5, 7-14, 7-26
 nonrouting, 1-4
 parameters,
 See Local node parameters
 Phase II, 1-4
 remote, 1-1
 routing, 1-4

- Nonrouting node, 1-5
- NSP, 1-3, 1-5, 2-1, 3-1, 7-2
- NTCC call,
 - See Connect Confirm Message
- NTCI call,
 - See connect Initiate Message
- NTCP call,
 - See Connect Reject Message
- NTDI call,
 - See Disconnect Message
- NTDM call,
 - See Network Data Message
- NTEV call,
 - See User Events
- NTIN call,
 - See Interrupt Message
- NTLA call,
 - See Link Abort Message
- NTLN call,
 - See Local node parameters
- NTLS call,
 - See Link Service Message

O

- Object
 - automatic startup,
 - See Automatic startup
 - identification,
 - See Receiver identity
 - multiple copies,
 - See Multiple copies of objects
 - type code,
 - See Receiver identity
- One-shot bit, 7-8

P

- Parameters,
 - bit values, 7-4
 - byte values, 3-2, 7-2
 - input, 3-2, 7-2
 - output, 7-2
 - with DEFINE or SET OBJECT, 6-4
- Password, 7-29, 7-35
- Path, 1-2
- Path length, 1-2
- Pending message queue, 4-6, 4-7, 5-3, 5-5, 5-11, 5-12, 6-13
- Phase II, 1-4
 - node, 1-4
- Phase III, 1-4
- Physical link, 1-1, 1-3, 1-5
- PICTURE statement, 3-2, 7-2
- Point-to-point lines, 1-5

- PPN, 7-22, 7-29, 7-35, 7-67, 7-68
 - default, 5-5, 6-6, 7-14
- Priority message,
 - See Interrupt Message
- Privileged job, 6-6
- Procedure Division, 3-2, 7-2
- Program,
 - local, 5-3, 5-4
 - network, 5-3, 5-4
 - remote, 1-1
 - source, 4-1
 - target, 4-1
- Program-regulated flow control, 4-8, 6-15
 - effect on DECnet/E programs, 6-16, 6-17
- message,
 - See Message flow control
- segment,
 - See Segment flow control
- selection, 5-7, 5-8, 6-15, 7-24, 7-32
- use of Link Service Message, 6-15, 6-17
- Project-programmer number, 1-3
 - See PPN
- Protocols, 1-4
 - DAP, 1-4
 - DDCMP, 1-3, 1-5
 - NICE, 7-17
 - NSP, 1-3
 - Transport, 1-3

Q

- Queue,
 - Data transmit, 6-10
 - Interrupt/Link Service, 6-10
 - management, 6-10
 - message limit, 6-11
 - pending message, 4-6, 4-7, 5-3, 5-5, 5-11, 5-12, 6-13

R

- Reason code, B-1
 - for Connect Reject, 5-10, 7-36
 - for Link Abort 5-12, 7-76, B-1, B-2
- Receive, 2-7, 5-3
 - selective, 5-11, 5-12
 - sleep, 7-58, 7-61
 - system call, 7-58
 - truncation of message,
 - Receive maximum, 5-7, 5-8, 5-9, 5-10, 7-24, 7-33
- Received formats,
 - Connect Confirm Message, 7-69
 - Connect Initiate Message, 7-65
 - Connect Reject Message, 7-70
 - Disconnect Message, 7-75

- Interrupt Message, 7-72
- Link Abort Message, 7-76
- Link Service Message, 7-73
- Local Data Message, 7-64
- Network Data Message, 7-71
- Receiver access,
 - See* Access
- Receiver ID Block,
 - See* RIB
- Receiver identity, 3-3, 4-7, 5-3, 5-7
 - consistency with access, 7-8
 - declaring, 6-1
 - name, 2-3, 3-3, 5-4, 6-1, 7-6, 7-21
 - object type code, 2-3, 3-3, 5-4, 6-1, 7-6, A-1
 - uniqueness, 6-2
- Receiver name,
 - See* Receiver identity
- Remote link address,
 - See* RLA
- Remote link status flags, 7-74
- Remote node, 1-1
- Remote program,
 - definition, 1-1
 - identification, 7-27
- Remove Receiver, 2-5, 5-2, 5-8
 - system call, 7-12
- RIB, 5-5, 6-2, 6-3, 6-13, 7-6
- RLA, 4-4
- Routing, 1-4
 - cost, 1-4
 - hop, 1-2
 - network diameter, 1-2
 - path, 1-2
 - path length, 1-2
- Routing node, 1-4
- RSTS/E system monitor, 1-4, 3-1, 5-1, 7-1

S

- Segment,
 - beginning of message, 2-5
 - definition, 6-15, 7-39
 - end of message, 6-16
 - maximum size, 6-15
 - request count, 7-39, 7-48, 7-49
- Segment flow control, 6-15
- Selective receives, 5-11, 5-12
- Send, 5-2
- SET EXECUTOR, 6-10
- SET OBJECT, 6-4
- Sleep,
 - in Declare Receiver, 7-8
 - in Receive, 5-3, 7-58, 7-61
- Sleep monitor call, 7-8
- Source program, 4-1
- Starting line number, 6-4

- Synchronous I/O, 2-4
- System buffer space, 4-7, 5-2, 5-4, 5-5, 7-9
- System calls,
 - general form with concise interface, 3-1
 - general form with extended interface, 7-1
 - summary, 5-1
- System library, 3-2, 7-1, 7-5
- System manager, 1-4, 1-5, 4-8, 6-1, 6-3, 6-4, 6-7, 6-8, 6-10, 6-11, 7-17

T

- Target program, 4-1
- Terminal Communication Utility,
 - See* TLK/LSN
- Timing,
 - automatic startup, 6-5
 - transmit queue management, 6-14
- TLK/LSN, 1-6, 2-3
- Transmit maximum, 5-2, 5-9, 5-10, 7-25, 7-33, 7-65, 7-69
- Transmit queue management, 6-10
 - effect on aborts, 6-11
 - effect on DECnet/E programs, 6-11
 - effect on disconnects, 6-11
 - message limit, 6-11
 - timing, 6-10
- Transport Protocol,
 - See* TRN
- TRN, 1-3, 1-6, 5-2
- Truncation of received messages, 5-3, 7-59

U

- UIC,
 - See* PPN 5-5, 5-27
- ULA, 4-2, 5-12
- User data accompanying messages, 5-3, 5-10, 5-12
- User events, 2-6, 5-2, 5-6
 - class and type, 7-17
 - entity, 7-18
 - parameter data, 7-18
 - system call, 7-17
- User Identification code,
 - See* PPN
- User link address,
 - See* ULA
- USING clause, 7-1

V

- VALUE IS clause, 3-2, 7-2
- Virtual terminal,
 - See* NET

READER'S COMMENTS

NOTE: This form is for document comments only. DIGITAL will use comments submitted on this form at the company's discretion. If you require a written reply and are eligible to receive one under Software Performance Report (SPR) service, submit your comments on an SPR form.

Did you find this manual understandable, usable, and well-organized? Please make suggestions for improvement.

Did you find errors in this manual? If so, specify the error and the page number.

Please indicate the type of user/reader that you most nearly represent.

- ☐ Assembly language programmer
- ☐ Higher-level language programmer
- ☐ Occasional programmer (experienced)
- ☐ User with little programming experience
- ☐ Student programmer
- ☐ Other (please specify) _____

Name _____ Date _____

Organization _____

Street _____

City _____ State _____ Zip Code _____

or
Country

digital




No Postage
Necessary
if Mailed in the
United States

BUSINESS REPLY MAIL

FIRST CLASS PERMIT NO.33 MAYNARD MASS.

POSTAGE WILL BE PAID BY ADDRESSEE

SOFTWARE DOCUMENTATION
1925 ANDOVER STREET TW/E07
TEWKSBURY, MASSACHUSETTS 01876



Do Not Tear - Fold Here and Tape

Cut Along Dotted Line